

YoLinux Tutorial: Linux System Administration



This tutorial covers managing, modifying and monitoring a Linux based computer.



- [Monitor the system.](#)
- [Process management.](#)
- [Filesystems and Storage Devices.](#)
- [Mounting CD's, Floppies etc.](#)
- [User Info.](#)
- [Creating a new system user account.](#)
- [Operating as root \(System Administrator\).](#)
- [RPM: Adding, updating software.](#)
- [RDIST: Remotely distributing and installing software/files](#)
- [CRON: Scheduling a task](#)
- [System log files.](#)
- [logrotate - Rotate log files.](#)
- [FIND: Using the find command](#)
- [Finding files on the system](#)
- [Managing Time](#)
- ["You have new mail."](#)
- [Text Terminal Configuration](#)
- [System Hardware Information.](#)
- [Red Hat automated system update: up2date.](#)
- [PERL Administration/Maintenanc](#)
- [Links](#)

Admin Notes:**YoLinux Admin Tutorials:**

- [init: The linux boot process.](#)
- [Quotas: Assign disk resource limits to users or groups.](#)
- [Managing groups and user security:](#)
- [Linux Networking tutorial](#)
- [Configuring an internet web server. Apache, DNS, FTP,...](#)
- [Internet Server Security](#)
- [System optimization:](#)
- [System re-configuration: Commands used to re-configure video, keyboard, mouse, date/time, sound card, NIC, Printer...](#)
- [Installing a new kernel RPM](#)

Monitoring the system:**Basic command line:**

pstree	Processes and parent-child relationships
top	Show top processes
ps -auxw	process status
vmstat	Monitor virtual memory
free	Display amount of free and used memory in the system. (Also: cat /proc/meminfo)
cat /proc/sys/vm/freepages	Display virtual memory "free pages". One may increase/decrease this limit: echo 300 400 500 > /proc/sys/vm/freepages
cat /proc/filesystems	Display filesystems currently in use.
uname -a	print system information
uptime	Tell how long the system has been running. Also number of users and system's load average.
w	Show who is logged on and what they are doing.
/sbin/lsmmod	List all currently loaded kernel modules. Same as cat /proc/modules
/sbin/runlevel	Displays the system's current runlevel.
hostname	Displays the system's name.

GUI/Graphical:

- [gkrellm](#) Graphical system monitor. (Additional package)
- [GLcpu](#) networked load meter. Very cool 3-D graphics.

List of tools:

- [tools for finding the status of your system](#)
-

Process Management:

The basic monitoring commands such as `ps`, `ps -auxw` and `top` will inform you of the processes running on your system. Sometimes a process must be terminated. To terminate a process:

1. Identify the process:
 - `ps -p`
OR
 - `ps -auxw`
OR
 - `top`
2. Kill the process:
 - `kill <process-id-number>`
 - `killall <command-name>`

This will perform an orderly shutdown of the process. If it hangs give a stronger signal with: `kill -9 <process-id-number>`. This method is not as sanitary and thus less preferred.

A signal may be given to the process. The program must be programmed to handle the given signal. See `/usr/include/bits/signum.h` for a full list. For example, to restart a process after updating it's configuration file, issue the command `kill -HUP <process-id-number>`

In the previous example, the **HUP** signal was sent to the process. The software was written to trap for the signal so that it could respond to it. If the software (command) is not written to respond to a particular signal, then the sending of the signal to the process is futile.

Identify all known signals: `fuser -l`

QPS: Also see the [GUI tool QPS](#). (Handles MOSIX cluster) This tool is outstanding for monitoring, adjusting nice values (priorities), issue signals to the process, view files the process is using, the memory, environment variables and sockets the process is using. RPM available from this site. It is so simple to use, no instructions are necessary. It can monitor a program to make sure it isn't doing something bad. I love this tool!!

Semaphores:

Note that some processes may use semaphores (shared memory) which may need to be cleaned up manually:

1. Identify the semaphores: [ipcs](#)
 - `ipcs -q` List share queues.
 - `ipcs -m` Shared memory.
 - `ipcs -s` List Semaphores.

- Remove the semaphores: `ipcrm -s <ipcs id>`

Example: If you are running Apache, you may see the following:

```
[root@node DIR]# ipcs -m
----- Shared Memory Segments -----
key          shmid      owner      perms      bytes      nattch     status
0x00000000   341504    nobody     600        46084      27        dest
```

ulimit:

Shell and process resources may be controlled and reported using the `ulimit` command. Display the limits of a shell using the bash command "`ulimit -a`". Limits can be set for the number of open files and processes, memory and virtual memory etc.

Filesystems and Storage Devices:

Hard Drive Info:

df -k	report filesystem disk space usage. (-k reports in Kbytes)
du -sh	Calculates file space usage for a given directory. (and everything uner it) (-s option sumerizes)
mount	Displays all mounted devices, their mountpoint, filesystem, and access. Used with command line arguments to mount file system.
showmount	Displays mount info for NFS filesystems.
<code>cat /proc/swaps</code>	Displays swap partition(s) size, type and quantity used.
<code>cat /proc/ide/hda/<i>any-file</i></code>	Displays disk information held by kernel.

Adding an extra hard drive:

- [fdisk](#) /dev/<drive> - **Enter the drive into the partition table**
- [mkfs](#) -t ext2 /dev/<drive> - **Create file system**
- [mount](#) -t ext2 /dev/<drive's device name> /<home2 or some suitable directory> - **Mount the drive**

Where the drive is /dev/hdb or some device as conforms to the Linux device names:

IDE drives are referred to as hda for the first drive, hdb for the second etc... IDE uses separate ribbon cables for primary and secondary drives. The partitions on each drive are referred numerically. The first partition on the first drive is referred to as hda1, the second as hda2, the third as hda3 etc ...

Linux IDE naming conventions:

Device	Description	Configuration
--------	-------------	---------------

/dev/hda	1st (Primary) IDE controller	Master
/dev/hdb	1st (Primary) IDE controller	Slave
/dev/hdc	2nd (Secondary) IDE controller	Master
/dev/hdd	2nd (Secondary) IDE controller	Slave

Note: SCSI disks are labeled /dev/sda, sdb, etc...

Use the command `cat /proc/partitions` to see full list of disks and partitions that your system can see.

See Linux devices:

- Kernel 2.4: (Red Hat 7.1+)
 - <file:///usr/src/linux-2.4/Documentation/devices.txt> - (local file)
- Kernel 2.2: (Red Hat 7.0-)
 - <file:///usr/src/linux/Documentation/devices.txt> (local file) - H. Peter Anvin

To make the drive a permanent member of your system and have it mount upon system boot, add it to your `/etc/fstab` file which holds all the file system information for your system. See [man page for fstab](#).

Example of existing `/etc/fstab` file:

```

/dev/sdb6          /                ext2    defaults    1 1
/dev/sdb1         /boot           ext2    defaults    1 2
/dev/cdrom        /mnt/cdrom     iso9660 noauto,user,ro 0 0
/dev/fd0          /mnt/floppy    auto    noauto,owner 0 0
none             /proc          proc    defaults    0 0
none             /dev/pts      devpts  gid=5,mode=620 0 0
/dev/sdb5        swap           swap    defaults    0 0

```

Add SCSI drive by adding line:

```

/dev/sdc1          /home2         ext2    defaults    1 2

```

At this point one may optionally check the file system created with the command: `fsck /dev/sdc1`

Note that `fsck` is NOT run against a mounted file system. Unmount it first if necessary. (`umount`) Also see the man page for:

- [cfdisk](#) - Curses based disk partition table manipulator. (very nice)
- [fsck](#)

[More on Mounting Filesystems and /etc/fstab](#)

Mounting other file systems:

- Mounting a CD: `mount -r -t iso9660 /dev/cdrom /mnt/cdrom`

This command should work for a Red Hat installation. Other distributions may require the following set-up:

```
ln -sf /dev/hdc /dev/scd0      My SCSI system
OR
ln -sf /dev/hdc /dev/cdrom    A more typical system

cd /mnt
mkdir cdrom
mount -t iso9660 -o ro /dev/cdrom /mnt/cdrom
```

Don't forget to un-mount the CD with `umount /mnt/cdrom`

[Potential Pitfall]: There is NO "N" in umount!!!

For trouble shooting your CD see your kernel documentation:

- Kernel 2.4: <file:/usr/src/linux-2.4/Documentation/ide.txt> (local file)
- Kernel 2.2: <file:/usr/src/linux/Documentation/ide.txt> (local file)

- [Mounting a Windows partition.](#)

- Mounting a floppy:

- Mount MS-DOS floppy: `mount /dev/fd1 -t vfat /mnt/floppy`

Also see the [YoLinux tutorial on using MS-DOS floppies with Linux.](#)

- Unix floppy: [See YoLinux Tutorial - Linux Recovery and Boot Disk Creation](#)

- Ramdisk: Using a portion of RAM memory to act like a superfast disk.

```
/sbin/mkfs -t ext2 /dev/ram
mount /dev/ram /mnt/ramd
```

[Potential Pitfall]: I've never actually tried this. Use at your own risk!

See [How to use a Ramdisk for Linux](#)

- [Mounting a Windows partition or other file systems.](#)
- [YoLinux Tutorial: File System Quotas](#)
- [Hard Disk Upgrade Mini How-To](#)

Increase open files limit:

```
cat /proc/sys/fs/file-max - Kernel configuration for max number of files
```

```
cat /proc/sys/fs/file-nr - Number of files presently open
echo 4096 > /proc/sys/fs/file-max - Set max file limit. (This is default)

cat /proc/sys/fs/inode-max - Kernel 2.2 configuration for max number of inodes
To change: echo 12288 > /proc/sys/kernel/inode-max
```

- See:
 - [proc](#) man page - process information pseudo-filesystem
 - Local file Kernel 2.2 (RH 7.0-): <file:///usr/src/linux/Documentation/proc.txt> (local file) - [\[WEB\]](#)

Reboot count and fsck:

After 20 reboots of the system, Linux will perform a file system check using `fsck`. This is annoying for systems with many file systems because they will all be checked at once. The individual file system's mount count may be changed so that they will be checked on a different reboot.

Check current reboot status: `/sbin/dumpe2fs /dev/sdb6 | grep '[mM]ount count'`

```
dumpe2fs 1.19, 13-Jul-2000 for EXT2 FS 0.5b, 95/08/09
Mount count:                2
Maximum mount count:        20
```

(This of course is specific for my system. Refer to your particular filesystem.)

Perform the previous command on all the filesystems to obtain their mount counts. Next change the mount counts for some of them.

```
umount /dev/sdb6
tune2fs -C 9 /dev/sdb6
mount /dev/sdb6
```

Now the filesystems will have an `fsck` performed on them on different system boots rather than all at the same time.

For home users who routinely shutdown and boot their systems, one can increase the maximum mount count: `tune2fs -c 40`

This feature can also be disabled: `tune2fs -c -1`

Check every week: `tune2fs -i 7`

System crash and disk check upon boot:

If the system crashes (due to power outage etc...) then upon boot the system will check if the disk was unmounted cleanly. If not you may get the following message:

```
Unexpected inconsistency; Run fsck Manually
...
*** An error occurred during the file system check.
```

```
*** Dropping you to a shell; the system will reboot
...
.
Give root password for maintenance
(or type Control-D for normal startup):
```

At this point enter the root password then run fsck:

```
(repair file system) 1# fsck -A -y
...
..
.
***** FILE SYSTEM WAS MODIFIED *****
...
..
.
(repair file system) 2# exit
```

The system will hopefully reboot properly at this point.

[fsck man page](#)

Note that `fsck` is NOT run against mounted file systems.

Journalized Filesystem EXT3:

Convert from ext2 to ext3:

- Red Hat 7.2 upgrade gives one the option to perform this file system upgrade.
- Manual method:
 - Convert: `tune2fs -j /dev/hda1`
 - Configuration file changes: `/etc/fstab` change ext2 to ext3

Also see:

- [Linux today: EXT3 info](#)
- [EXT3 FAQ](#)

Other journalized file systems: SGI XFS, IBM JFS and reiserfs. For files larger than 2Gb use [SGI XFS and the SGI Linux Red Hat RPM or Red Hat ISO CD install image](#).

User Info:

Commands:

- [who](#) Displays currently logged in users.
Use `who -uH` for idle time and terminal info.
- [users](#) Show all users logged in.
- [w](#) Displays currently logged in users and processes they are running.
- [whoami](#) Displays user id.
- [groups](#) Display groups you are part of.

- Use `groups user-id` to display groups for a given user.
 - [set](#) Display all environment variables in your current environment.
 - [id](#) Display user and all group ids.
Use `id user-id` to display info for another user id.
 - [last](#) Show listing of last logged in users.
 - [history](#) Shell command to display previously entered commands.
-

Creating a new system user account:

- **Command Line Method:** (My preference)
 - [useradd](#) : Add a user to the system.
 - `useradd -G floppy` : Will grant the user read/write privileges to the floppy (/dev/fd0) upon creation of user by adding user to group floppy.
 - [userdel](#) : Delete user from system. Purges user from /etc/passwd, group and shadow files
 - `userdel -r` : Delete user and remove his home directory from the system. Other files will remain.
 - [passwd](#) : Assign a password to the user

Also see man page for: [usermod](#).

Configuration file for useradd command: /etc/default/useradd

Default directory configuration and files for a new user are copied from the directory /etc/skel/. The default shell is called bash (bsh) and is a cross of the UNIX ksh and csh command shells. The users personal bash shell customizations are held in \$HOME/.bashrc.

- **GUI Method:** - Using system administration tool linuxconf:
 - Start linuxconf:
 - RH 5.2: Select Start + Programs + Administration + linuxconf .
 - RH 6+: Select Gnome Start icon (located lower left corner) + System + Linuxconf .
 - Add the user: Select options Config + User accounts +Normal + User accounts + select button Add . There is also the option of adding the user to additional groups. (I.e enter floppy under the heading Supplementary groups andthen Accept) For a list of groups, the group names should be separated by a simple space. This tool will allow you to set default directories, shells, add rules about passwords, set e-mail aliases, group membership and disk quotas. One can modify or delete users from linuxconf as well.
 - Set user password: After creating the user, use options Config + User accounts + Normal + User accounts .Select the user from the list. Then select the Passwd button. This will allow you to enter an initial password for the account.
- **File Editing Method:** - (as root) Edit files to add/remove a user
 - Create user entry in /etc/passwd
`user:x:505:505:Mr. Dude User:/home/user:/bin/bash`
 - Create group: /etc/group
`user:x:505:`
 - Create home directory:
`cd /home`
`mkdir user`

- Copy default files:


```
cp -pR /etc/skel/. /home/user
chown -R user:user /home/user
```
- The creation of `/etc/shadow` and `/etc/gshadow` require the execution of a program to encrypt passwords. Use the commands [pwconv](#) and [grpconv](#) to synchronize the shadow files.
- Assign a password: `passwd user`
- Also see:
 - Shadow integrity verification: `grpck [-r] [group shadow]`
 - File editor: [vipw](#).

Operating as root: (System Administrator)

If you are planning to administer the system, one would login as root to perform the tasks. In many instances one would be logged in as a user and wish to perform some "root" sys-admin tasks. Here is how:

- Switch user account to root: Shell/command mode:


```
su -
```

 and then you will be prompted to enter the root password.
- GUI mode:
 - Allow console to be accessed by another user from same system:


```
xhost +localhost
```
 - Switch user and then you will be prompted to enter the root password:


```
su -
```
 - Set environment variable for GUI sessions so that display used is local display. (Remember X-windows is a network enables windowing system)


```
export DISPLAY=:0.0
```
 - Launch GUI application.

Note: The `su` command was issued with a trailing "-" which means that the login process will follow the login steps of the user and will execute its profiles. Since no user id was given, the default is root user. This all would be equivalent to: `su - root`

Some systems may be configured so that only the switch user (`su`) command may be required without all of the X-window configuration.

RPM - Redhat Package Manager:

The `rpm` command is used to manage software applications and system modules.

RPM Command	Description
<code>rpm -qip program_package-ver.rpm</code>	Query for information on package and list destination of files to be installed by the package.

<code>rpm -Uvh program_package-ver.rpm</code>	Upgrade the system with the RPM package
<code>rpm -ivh program_package-ver.rpm</code>	New Install
<code>rpm -Fvh program_package-ver.rpm</code>	Freshen install. Removes all files of older version.
<code>rpm -q program_package</code>	Query system RPM database (/var/lib/rpm), to see if package is installed.
<code>rpm -qi program_package</code>	Query system RPM database for info/description on package (if installed)
<code>rpm -ql program_package</code>	List all files on the system associated with the package.
<code>rpm -qf file</code>	Identify the package to which this file belongs.
<code>rpm -e program_package</code>	Uninstall package from your system
<code>rpm -qa</code>	List ALL packages on your system. Use this with grep to find families of packages.
<code>rpm -K --nogpg *.rpm</code>	Non sure if RPM downloded ok? Verify md5 sum.

RPM Flag	Description
<code>--nodeps</code>	RPM flag to force install even if dependancy requirements are not met.
<code>--force</code>	Overwrite of other packages allowed.
<code>--notriggers</code>	Don't execute scripts which are triggered by the installation of this package.
<code>--root /directory-name</code>	Use the system chrooted at /directory-name. This means the database will be read or modified under /directory-name. (Used by developers to maintain multiple environments)
<code>--ignorearch</code>	Allow installation even if the architectures of the binary RPM and host don't match. This is often required for RPM's which were assembled incorrectly

Notes:

- Many times, (like with glibc library or Netscape RPMs etc) it is necessary to mention two or more packages on the command line. The rpm command will account for the co-dependancy of the packages. i.e.:

```
rpm -ivh abc-package1-i386.rpm abc-package2-i386.rpm abc-package3-i386.rpm
```
- Configuration information is stored in `/var/lib/rpm`
- When installing additional RPM's from the Red Hat CD, `cd` to the RPMS directory on the CD which contains the packages to be installed.

Usefull man pages:

- [rpm](#)

- [glint](#) - RH 5.2 GUI
- [gnorpm](#) - RH 6.0+ GUI

Also see:

- [RPM HowTo](#).
- [RMP.org Home Page](#)
- [RpmLevel](#) - Utility to manage RPM levels on your system.
- [AutoRPM](#) - Mirror RPMs from an FTP site, keep installed RPMs consistent with an FTP site or local directory
- [Alien](#) - package converter between rpm, dpkg, stampede slp, and slackware tgz file formats.

RDIST: Remotely distributing and installing software/files

The command `rdist` helps the system administrator install software or update files accross many machines. The process is launched from one computer.

Command: `rdist -f instruction-file`

Instruction file:

```
files=(  
  /fully-qualified-path-and-file-name  
  /next-fully-qualified-path-and-file-name  
)  
dest = ( computer-node-name )  
  
${files} -> ${dest}  
  
install /fully-qualified-directory-name-of-destination;
```

Mote: Dont forget the ";" at the end.

For more info see the [rdist man page](#)

Also see the [rsync man page](#) to migrate file changes.

CRON - Scheduling a task:

Add shell scrip to have run hourly, daily, weekly or monthly into the appropriate directory:

- /etc/cron.hourly/
- /etc/cron.daily/
- /etc/cron.weekly/
- /etc/cron.monthly/

These are preconfigured schedules. To assign a very specific schedule add a line to the `/etc/crontab` file. Cron entries may also be added to a crontab formatted file located in the directory `/var/spool/cron/`.

Scheduling access and control:

- The administrator can allow users to use this facility with specific control by using the `/etc/cron.deny` and `/etc/cron.allow` files.
- The `at` facility may be controlled with the `/etc/at.deny` and `/etc/at.allow` files.

Man pages:

- [cron](#)
- [crontab](#)
- [at](#)

System log files:

```
/var/log/messages - system messages
    /secure       - Logging by PAM of network access attempts
    /dmesg        - Log of system boot. Also see command dmesg
    /boot.log     - Log of system init process
    /xferlog.1    - File transfer log
    /lastlog      - Requires the use of the lastlog command to examine content
    /maillog      - log from sendmail daemon
```

Note: The [lastlog](#) command prints time stamp of the last login of system users. (Interprets file: `/var/log/lastlog`)

logrotate - Rotate log files:

Many system and server application programs such as Apache, generate log files. If left unchecked they would grow large enough to burden the system and application. The `logrotate` program will periodically backup the log file by renaming it. The program will also allow the system administrator to set the limit for the number of logs or their size. There is also the option to compress the backed up files.

Configuration file: `/etc/logrotate.conf`

Directory for logrotate configuration scripts: `/etc/logrotate.d/`

Example logrotate configuration script: `/etc/logrotate.d/process-name`

```
/var/log/process-name.log {
    rotate 12
    monthly
    errors root@localhost
    missingok
    postrotate
```

```

        /usr/bin/killall -HUP process-name 2> /dev/null || true
    endscript
}

```

The configuration file lists the log file to be rotated, the process [kill](#) command to momentarily shut down and restart the process, and some configuration parameters listed in the [logrotate man page](#).

Using the find command:

[Find man page](#)

Form of command: `find path operators`

Examples:

- Search and list all files from current directory and down for the string *ABC*:
`find ./ -name "*" -exec grep -H ABC {} \;`
- Find all files of a given type from current directory on down:
`find ./ -name "*.conf" -print`
- Find all user files larger than 5Mb:
`find /home -size +5000000c -print`
- Find all files owned by a user (defined by user id number. see `/etc/passwd`) on the system: (could take a very long time)
`find / -user 501 -print`
- Find all suid and setgid executables:
`find / \(-perm -4000 -o -perm -2000 \) -type f -exec ls -ldb {} \;`
`find / -type f -perm +6000 -ls`

Note: suid executable binaries are programs which switch to root privileges to perform their tasks. These are created by applying a "sticky" bit: `chmod +s`. These programs should be watched as they are often the first point of entry for hackers. Thus it is prudent to run this command and remove the "sticky" bits from executables which either won't be used or are not required by users. `chmod -s filename`

- Find all world writable directories:
`find / -perm -0002 -type d -print`
- Find all world writable files:
`find / -perm -0002 -type f -print`
`find / -perm -2 ! -type l -ls`
- Find files with no user:
`find / -nouser -o -nogroup -print`
- Find files modified in the last two days:
`find / -mtime 2 -o -ctime 2`

Also see:

- [gFind](#) - GUI front-end to the GNU find utility

Finding/Locating files:

- [locate/slocate](#) Find location/list of files which contain a given partial name
- [which](#) Find executable file location of command given. Command must be in path.
- [whereis](#) Find executable file location of command given and related files
- [rpm](#) -qf *file* Display name of RPM package from which the file was installed.

Note: The script `/etc/cron.daily/updatedb.cron` generates the index for the **locate** command. It will generate the database `/var/lib/locatedb`.

Managing Time:

Set System Time:

- Print the time returned by the remote host: `rdate -p hostname`
- Set the system time to the returned time: `rdate -s hostname`

Try: `rdate -p time.ucla.edu`

[See List of public NTP Time Servers](#)

Note: Typically many web servers set their time to GMT due to the world wide nature of their service. Internally UNIX systems use Coordinated Universal Time (UTC) which is the number of seconds since Jan 1, 1970 0 hrs. "Calendar Time" is then calculated based on your time zone and whether you are on Standard or Daylight Savings time.

Sync System Time:

The `timed` (time server daemon) allows one to synchronize the host's time with the time of another host. This is a master - slave configuration. See the [timed](#) and [timedc](#) man pages.

[Freshmeat list of time daemon software](#)

Shell environment variables:

- TZ: Time Zone
 - i.e.: `export TZ=PST8PDT`
 - Zone format: `[time standard] [offset] [dst]`
 - See directory: `/usr/share/zoneinfo/`
 - File: `/usr/share/zoneinfo/PST8PDT` not viewable.
 - See file: `/usr/share/zoneinfo/zone.tab`
 - See [tzset](#), [tzselect](#) man pages.
- TIMEFORMAT:
 - For more info see the man page for your shell. (i.e. [man bash](#))

Network Time Protocol (NTP):

The daemon `ntpd` will continually monitor time. The init script `/etc/rc.d/init.d/ntpd` issues the command `/usr/sbin/ntpdate` to set the time. It expects time servers to be listed in `/etc/ntp/step-tickers`.

You may also issue the command: `/usr/sbin/ntpdate -b -s time.nist.gov ns.arc.nasa.gov tick.usno.navy.mil`

This will hang without an internet connection!!!

Time servers:

- `time.nist.gov`
- `ns.arc.nasa.gov`
- `tick.usno.navy.mil`

Links:

- [Network Time Protocol \(NTP\)](#)

"You have new mail.":

When you login, this message may greet you. The system will often send a mail message to the "root" user after the completion of some cron jobs, software installation or as an error message meant to alert the system administrator. Type the console command "mail". The following simple commands will help you navigate through this simple mail client.

Mail command	Description
?	List commands (Help)
h	Print mail headers
h 1	Print mail headers starting with message number 1
q	Quit
t	Type current message
t 4	Type out message 4 to the console
n	Type out next message
d	Delete the active message
d 4	Delete message number 4

- When typing a "cntl-m" is just like hitting the "Enter" key. If you want to enter the "cntl-m" as part of the entry to the `stty` command then prefix it with "cntl-v" so that the "cntl-m" "escaped" from acting as a terminal directive but instead acts as command input.
 - Check terminal type: `echo $TERM`
 - Set terminal type: `export TERM=xterm`
This is a very common fix for many remote terminal problems.
-

Hardware Info:

[lsdev](#) List devices and info on system hardware. Also IRQ's.

[/sbin/lspci](#) list all PCI devices (result of probe) Also `lspci -vvx` and `cat /proc/pci`

`cat /proc/interrupts` List IRQ's used by system.

`cat /proc/ioports` List I/O ports used by system.

`cat /proc/dma` List DMA channels and device used by system.

`cat /proc/cpuinfo` List info about CPU.

Also See:

- Kernel 2.4: (Red Hat 7.1+)
 - <file:///usr/src/linux-2.4/Documentation/devices.txt> - (local file)
 - Kernel 2.2: (Red Hat 7.0-)
 - See Linux devices: <file:///usr/src/linux/Documentation/devices.txt> (local file) - H. Peter Anvin
 - Local file Kernel 2.2: <file:///usr/src/linux/Documentation/proc.txt> (local file) - [\[WEB\]](#)
-

PERL Administration/Maintenance:

At some point you will be required to administer the installation of PERL modules.
Installation can be done:

- Manually:
 - Un-zip/Un-tar module: `tar xzf yourmodule.tar.gz`
 - Build with PERL makefile:
 - `perl Makefile.PL`
 - `make`
 - Install: `make install`

See: [CPAN module install page](#)

- Automatically: (preferred)

```
# perl -MCPAN -e shell          - First time through it will ask a bunch of questions. Answer "no" t
...
..
```

```

cpan> install URI
...
..

cpan> i /PerlMagick/           - Inquire about module. (Search by keyword)
Distribution   J/JC/JCRISTY/PerlMagick-5.36.tar.gz
Module        Image::Magick   (J/JC/JCRISTY/PerlMagick-5.36.tar.gz)

cpan> install Image::Magick
...

cpan> install Image::Info
...

cpan> install IO::String
IO::String is up to date.

cpan> help

```

This method rocks! It connects to a CPAN server and ftp's a gzipped tarball and installs it. First time through it will ask a bunch of questions. (Answer "no" to the first question for autoconfigure.) Defaults were good for me. The only reason to manually configure this is if you are using a proxy. It then asks for your location (i.e. North America) and country. I entered a number for the first CPAN server but after that the actual URL was cut and pasted in whole.

If it fails, you must load the appropriate RPMs and retry using "`force install module-name`"

```
PERL update: perl -MCPAN -e 'install Perl'
```

Automated System Updates: (Red Hat 7.1)

Execute the following commands (in order given) to perform an automatic system update:

1. `/usr/bin/rhn_register` :You must first register your system with the Red Hat database. This command will perform a hardware inventory and reporting of your system so that Red Hat knows which software to load to match your needs.
2. `/usr/bin/up2date-config` :This allows you to configure the "up2date" process. It allows you to define directories to use, actions to take (i.e. download updates, install or not install, keep RPM's after install or not), network access (i.e. proxy configuration), use of GPG for package verification, packages or files to skip, etc.
3. `/usr/sbin/up2date` :This command will perform an audit of RPM's on your system and discover what needs to be updated. It gives you a chance to unselect packages targeted for upgrade. It will download RPM packages needed, resolve dependencies and perform a system update if requested.

[Potential Pitfall]: This works quite well but it is not perfect. Red Hat 7.1 Apache upgrade to 1.3.22 changed the configuration completely. (Beware. manual clean-up and re-configuration is required). When `up2date` finds the first messed up dependency it stops to tells you. You then have to unselect the package. It then starts again from the beginning.

Also see:

- `up2date-gnome`
- `rhn_register-gnome`

The default download directory for `up2date` is `/var/spool/up2date/`. RPM packages and support files are downloaded to this directory.

An automated alert to the need to update utilizes the `rhnsd` which can be started by issuing the command: `/etc/rc.d/init.d/rhnsd start`

[Red Hat Network: User Reference Guide](#)

Links:

- [Linux Sys Admin guide](#) - by Lars Wirzenius, Joanna Oja
- [LINUX Rute Users Tutorial and Exposition](#) - Paul Sheer
- [Process Monitoring HowTo](#) - Alavoor Vasudevan
- [Basic commands](#) - UCLALUG Guide
- [TrinityOS: Guide to Configuring Your Linux Server for Performance, Security, and Managability](#)
- [LAME: Linux Administration Made Easy](#) by Steve Frampton
- [KALUG: Linux Administrators FAQ List](#)
- [Caldera System Administrator's Guide](#)
- [List of commands](#)
- [SysAdmin Magazine](#) - Journal for Unix System Administrators
- [LinuxConf](#) (Solucorp)
- Shell Script Resources:
 - [Bash: Linux terminal command guide](#)
 - [SHELLdorado](#) - UNIX shell scripting resource
 - <http://theory.uwinnipeg.ca/UNIXhelp/srpt/index.html>

SysAdmin Tools:


- [Webmin](#)

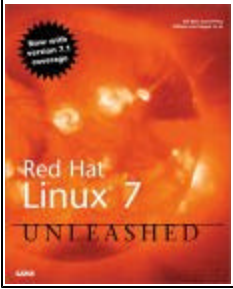

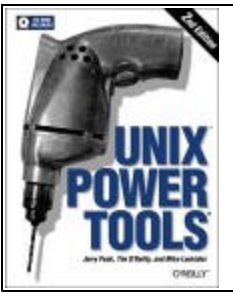


Alternate configurations:

- [Single-User Booting Under Linux](#)
- [Diskless-HOWTO](#)
- [Diskless-root-NFS-HOWTO](#)



Books:

	<p>"Red Hat Linux 7 Unleashed" by Bill Ball, David Pitts, William Ball Sams, ISBN# 0672319853</p> <p>I have the Red Hat 6 version and I have found it to be very helpful. I have found it to be way more complete</p>	
--	---	---

	<p>than the other Linux books.</p>	
	<p>"UNIX Power Tools" by Jerry D. Peek, Tim O'Reilly, Mike Loukides O'Reilly & Associates, ISBN# 1565922603</p> <p>This book is a requirement for any serious system administrator or developer. One of my favorite UNIX books.</p>	 

Return to <http://YoLinux.com> for more Linux links, information and tutorials
 Return to [YoLinux Tutorial Index](#)

REBUILDING YOUR CREDIT? [Click Here](#)

Apply Safely Online - No Security Deposit 

Copyright © 2000, 2001, 2002 by
 Greg Ippolito