HOWTO: How to stay updated

Table of Contents

HOWTO: How to stay updated

Kjetil Bakkeskaug, Kjell Sundby and Stein Gjoen, `sgjoen@nyx.net`

v0.32, 20 May 2002

This document describes how to stay updated and abreast of the development that takes place in the Linux world of development. Although most of this text is Linux specific there is also a lot of general information on searching efficiently for specific information that can be useful for a wider audience.

---

1. Introduction

Development in the world of Linux takes place at an incredible speed and it can be difficult to keep abreast with the latest development. This HOWTO gives you a few guidelines on how to get the information you need, fast and efficiently. Most are quite familiar with using the World Wide Web (WWW) and Usenet News but as will be shown here there are many other methods that can be as good or even better.

There are now many new translations available and special thanks go to the translators for the job and the input they have given:

- Polish translation by Tomasz Sienicki `tsca (at) cryogen.com`
- French Translation by Jean−Albert Ferrez `Jean−Albert.Ferrez (at) epfl.ch`
- Portuguese Translation by Duarte Loreto `dnloreto (at) esoterica.pt`
- Japanese translation by Satoru Takahashi `hisai (at) din.or.jp`

1.1 Copyright

If you have any questions, please contact <{linux−howto@metalab.unc.edu}>

1.2 Disclaimer

Use the information in this document at your own risk. We disavow any potential liability for the contents of this document. Use of the concepts, examples, and/or other content of this document is entirely at your own risk.

All copyrights are owned by their owners, unless specifically noted otherwise. Use of a term in this document should not be regarded as affecting the validity of any trademark or service mark.

Naming of particular products or brands should not be seen as endorsements.

You are strongly recommended to take a backup of your system before major installation and backups at regular intervals.

1.3 News

Updated links to LinuxNetMag

Did a major link check but linkrot sets in quickly.

Added link to Penguin Magazine.

Added more information on the workings of mailing lists. Also proper indexing is now added.

Renamed Dejanews to Deja and one chapter title. Also added note on translation now underway. Minor typos fixed.

Added link to Polish translation

Added link to Linux Focus and Linux Magazine.

Added links to French and Japanese translations as well as several online resources.

Link rot is setting in, numerous corrections made. Also expanded on searching choices.

1.4 Credits

```
Tomasz Sienicki | tsca <tsca (at) cryogen.com>
Satoru Takahashi <hisai (at) din.or.jp>
```

---

2. Documentation Installed on Your Hard Disk

When you start out with a Linux installation you will normally get quite a bit of information along, not just the installation pamphlet but also substantial online help and information files as well as HOWTO files. This gives you a good starting point but after a while you will find yourself interested in knowing more, updating your system and basically staying informed. For simplicity this kind of information is here divided into several types, the type you subscribe to, information you search for as well as a bit on getting more specific help efficiently.

Even if you don't get printed information of some kind with your Linux packages you will along with any self respecting distribution get a number of directories with documentation of some kind, ranging from the tersest *README* files for most software packages to the more in depth *HOWTOs*, of which this is one.

Have a look in the document archive where most packages store their main documentation and README files etc. Also you will here find the HOWTO archive of ready formatted HOWTOs and also the mini−HOWTO archive of plain text documents.

The kernel source is, of course, the ultimate documentation. In other words, *use the source, Luke*. It should also be pointed out that the kernel comes not only with source code which is even commented (well, partially at least) but also an informative documentation directory. If you are about to ask any questions about the kernel you should read this first, it will save you and many others a lot of time and possibly embarrassment.

The online documentation is excellent for browsing and searching but don't dismiss the printed version altogether; if you cannot even get the machine to boot, how are you going to be able to read that piece of information you need to get the system going again?

---

3. Subscription

This basically means you set up a subscription of some sort and then follow the news as they come in. Be careful not to bite over more than you can chew, *information overload* is more than a buzzword. There is also rather more junk out there than is just annoying, it is a real problem these days. Read critically and be prepared to unsubscribe.

There are two distinct medias for getting continuous updates: news and mailing lists, though sometimes news is gatewayed to mail and vice versa. In general news is a larger volume, larger noise source compared to mailing lists. Trying to follow too many newsgroups is like drinking from a fire hose.

3.1 Usenet News

Getting access to Usenet News is outside the scope of this HOWTO, there are others that will help you with getting and reading News directly from a Linux system. If you have never used News before you should be careful to read the introductory information thoroughly. In spite of looking like an anarchy it does have its own distinct culture, follow a newsgroup for some time before posting yourself. Most importantly, look out for postings called *Frequently Asked Questions* or *FAQ* as they will show you the ropes for the group it is posted to, and most likely give you the answer to what you are looking for. Asking an FAQ will earn you severe negative credibility points as well as a place in many killfiles.

FAQs should be posted regularly but if you cannot find it you can always find it at the main FAQ archive at MIT.

These are also available as web pages.

Still, there is a lot of noise, spam and junk in News and this is where killfiles come in. You will need a news reader with killfile capability and when properly set up it will scan through a newsgroup according to a search key of your own design and mark all flagged postings as already read so you don't have to be bothered by the noise. This improves the signal–to–noise ratio and lets you concentrate on the important parts. Similarly, if you make noise in News others will killfile you so if you later were to ask for help they will never see your post.

Now to business: the following is a list of useful newsgroups:

- announcements
- answers, HOWTOs, FAQs etc.
- development of applications
- development of the system
- hardware
- misc
- networking
- setting up linux
- X11 on linux

Many national hierarchies also have Linux groups, such as the Norwegian Linux groups. If you cannot find your national or local group you might be able to use Deja to find the names for you.

3.2 Mailing Lists

Unlike Usenet News a mailing list is centralised, someone sends a mail to the server and the server in return mails everyone that is subscribed to that particular list. These lists are generally low volume but also very low noise. Any breaches of the charter will be looked harshly upon. Equally seriously it will delay the development or the project that the list is dedicated to. When you subscribe you will normally get an introductory mail describing the charter, again you are strongly recommended to read this very carefully.

There are many types of mail servers that can handle a list and you will need some information on how and where you can subscribe.

One of the most common list servers is *Majordomo* which is what the list server at vger.rutgers.edu is running. To learn how it works you send a mail message with the word `help` in the body. If you send it something it cannot parse you will get this help message anyway. If you instead mail it the word `lists` you will be returned a list of all mailing lists it serves, and that can be a considerable number.

Other mailing lists use several addresses, one where you send your requests such as `subscribe` and `unsubscribe`, and one where you send your contributions to the list which is usually also the address from which the list is also redistributed to you. Again, sending it the message `help` or something it cannot parse will give you the help information. An example: you send the word `subscribe` to the address `corned-beef-requests@somelistserver.org` and then you get mail from and contribute to the list address `corned-beef-list@somelistserver.org` until you unsubscribe.

A few tips before you start sending in to mailing lists:

- Do not send `subscribe` etc. to the list itself, only to the server address, otherwise you will look silly and you will annoy people. There can be several thousand subscribers to a list and if such errors were to pour in the noise would be too much.
- When you subscribe you will often get an introductory message sent to you automatically. Read it carefully as this should answer most of the initial questions.
- Do not gateway mailing lists to news without asking first as this can cause mailing loops as well as spam.

As mentioned above, vger.rutgers.edu. is one of the main mailing list servers and here is an abbreviated index of what is available for the Linux community:

- linux−8086 (Linux on Intel 8086 processors)
- linux−admin (Administration of Linux systems)
- linux−alpha (Linux on the Alpha processor platform)
- linux−apps (Applications)
- linux−arm (Linux on the Arm processor platform)
- linux−bbs (Linux Bulletin Board Systems)
- linux−c−programming (C−programming with Linux)
- linux−config (Configuration)
- linux−console (Console)
- linux−diald (Dial on demand daemon)
- linux−doc (Linux documentation)
- linux−fido (Linux fido network)
- linux−fsf (Linux and the Free Software Foundation)
- linux−ftp (Linux File Transfer Protocol)
- linux−gcc (Linux and the GNU C compiler)
- linux−gcc−digest (Digests of the above)
- linux−hams Amateur Radio and Linux discussions
- linux−hppa (Linux on the HP Precision Architecture processor platform)
- linux−ibcs2 (Linux and the Intel Binary Compatibility system)
- linux−ipx (Linux and Novell IPX networking protocol)
- linux−isdn (Linux and Integrated Services Digital Network)
- linux−japanese (Linux and Japanese extensions)
- linux−kernel (Linux kernel)
- linux−kernel−announce (Announcements for the above)
- linux−kernel−digest (Digests of the linux−kernel list)
- linux−kernel−patch (Linux kernel patches)
- linux−laptop (Linux on laptops)
- linux−linuxss
- linux−lugnuts (Linux User Groups)
- linux−mca (Linux and the IBM Micro Channel Architecture bus)
- linux−mips (Linux on the MIPS processor platform)
- linux−msdos (Linux and MSDOS)
- linux−msdos−devel (Linux − MSDOS development)

- linux−msdos−digest (Digest of the linux−msdos list)
- linux−net (Linux and networking)
- linux−new−lists (New mailing lists for Linux)
- linux−newbie (Linux and the inexperienced)
- linux−newbiew
- linux−nys
- linux−oasg
- linux−oi
- linux−opengl (Linux and the OpenGL graphics system)
- linux−pkg
- linux−ppp (Linux and the Point−to−Point Protocol)
- linux−pro
- linux−qag
- linux−raid (Linux and Redundant Array of Inexpensive Drives)
- linux−scsi (Linux and Small Computer Systems Interface)
- linux−serial (Linux and the serial system)
- linux−seyon (Linux terminal system)
- linux−smp (Linux Symmetric Multi Processing)
- linux−sound
- linux−standards
- linux−svgalib (Linux and the SVGA library)
- linux−tape (Linux and tape storage)
- linux−term (A Linux communications program)
- linux−userfs (Linux User File System)
- linux−word
- linux−x11 (Linux and the X Window System, Version 11)
- linux−x25 (Linux and the X25 Networking Protocol)
- sparclinux (Linux on the SPARC processor platform)
- ultralinux (Linux on the Ultra−SPARC processor platform)

There are of course a number of other lists on other server. As this is in a constant state of flux there is little point in naming all but the most important here. Instead you could check out a web page that maintains such a list of lists on various servers of interest to Linux users. It also offers an user friendly interface to subscribe or unsubscribe to the various lists directly.

There is also a web page listing a huge number of lists concerning much more than Linux at Publicly Available Mailing Lists.

3.3 Magazines

Many have been disappointed at the lack of information on Linux in the trade press. This is probably because certain commercial products would not stand up for any comparison and the advertisers would not stand it at all. Fortunately there is one Linux specific journal, called the *Linux Journal*. More information on subscription etc. can be found at the SSC home page. A table of contents is usually also available online.

Another commercial paper magazine is the Linux Magazine which also offers table of contents and some excerpts online.

Also Unix Review (formerly known as Performance Computing) gives a lot of Linux coverage.

Some popular e−zines are

- Linux Gazette
- Linux Focus (available in a number of languages)
- Linux Research
- Linux Today
- Linux News
- Linux Netmag (English) and Linux Netmag (German)
- Penguin Magazine

and probably a few others as new ones seem to pop quite frequently.

Check out LinuxHQ for up to date information on current news services.

New web pages with literally daily news on linux are popping up everywhere, many are quite professional in layout as well as in scope. One of the bigger ones is Freshmeat which serves out news daily.

For those who cannot afford the time to follow the net on an hourly basis yet need the important news quickly there is the Linux Weekly News, which gives you a weekly update of important news, including securities alerts and also announcements of new and updated software.

You can also find directions to IRC online chat lines at Linux.com.

There are also a number of more hardware oriented web sites worth visiting, such as Toms Hardware, Anandtech for general hardware reviews, and Storage review for the latest in disk, tape and other storage technology..

---

4. Meetings

Linux has been created through a massive networked effort, mostly by heavy use of the Internet. Still, there is the chance of meeting real people, face to face, in Linux user groups (LUG) that are all over the world. Search the lists that are published regularly, there could be one near you.

Conferences, install fests, creating new user groups and more is regularly announced on Usenet News announcements. Such events are excellent venues for staying on top of events and also for getting help.

---

5. Searching

There are many avenues open when searching for something particular. Remember you can also use the web search engines and that some, like

- Altavista
- Excite
- Hotbot

can also search Usenet news.

There are numerous search engines available but they are not all equal neither in method of rating relevance nor in size of database. For established, authoritative pages I recommend using Google since it scores pages by the number of links pointing to them. People often link to pages they find important or useful and Google uses this.

Google takes time to score a page properly so it is not that useful for searching for the newest or more exotic topics. For that I normally recommend either the very fast and wide spanning AllTheWeb or meta engines, web systems that forwards your request to a number of other search engines, receives the results, collates and scores before presenting you with the result. My favourite meta engine is Go2Net formerly known as MetaCrawler.

Such preferences are all subjective and the systems improve continuously so you need to experiment yourself. There is no perfect search engine and as less than 10 percent of all pages are indexed you need to try several engines if you don't succeed at first.

Also remember that Deja is a dedicated news searcher that keeps a news spool from early 1995 and onwards.

Even though more and more things take place of the web these days, do not forget that there is a lot of information available on the various ftp servers around the world. Some web search engines also index ftp servers but the tool of choice is still the Archie servers, systems that regularly scan major ftp servers around the world and keep lists of files. These can be accessed in many ways, either by Archie clients like `archie` or the X11 version `xarchie` which should be available on any well maintained linux system. Failing that you can access archie servers using telnet to any of the servers listed

- Australia
- Austria
- Belgium
- Finland
- Germany
- Korea
- Italy
- Japan
- Poland
- Sweden
- Spain
- United Kingdom
- United States
- United States

Of course you should try to use the server closest to you, and to see the list of current server you can either start the Archie client with no arguments or, if telnetting, by querying the server. Online help is available. Unfortunately not all servers are synchronised, so you might have to search a few before finding what you are looking for.

Recently a more user friendly ftp index server entered the net, the ftpsearch engine, featuring many options and with a rather stark interface, in the best Unix tradition.

Most of these offer help on efficient searching techniques, reading this can speed up your searches enormously. Investing a little time here will pay off in the long run.

If you have trouble getting onto the Internet but have mail then you should have a look at the *access via mail* FAQ. Naturally you can get it over e−mail using the US, Canada and South America server entering only this line in the BODY of the note:

```
send usenet/news.answers/internet-services/access-via-email
```

or Europe, Asia etc server entering only this line in the BODY of the note:

```
send lis-iis e-access-inet.txt
```

or look through your news spool if you have one locally in news.answers newsgroup.

Finally, you might wish to get more information from a person, say an author of a software package. usually you can find the e–mail address in the accompanying documentation which normally would be in the documentation subdirectory but failing that and also if the given address is no longer valid you could find help in the FAQ for finding e–mail addresses.

6. Conclusion

Finding information fast and efficiently is more of an art than a science and we still have not touched on the really difficult part: how do you determine the actual *quality* of the information? It is outside the scope of this HOWTO to tell you that but it is still something you should keep in mind. You should at least check the information is recent enough to be current to your problem.

As a bare minimum you should ensure a minimum of validity of the documentation to avoid misleading or malicious advice. A surprising number of people suggests things like `rm -rf /` as a solution for a given problem. Some see it as an obvious prank, the unaware can end up destroying his or her setup. Just to avoid such things you should check out a few things before rushing ahead:

- Is there a name attached to the document? If people are serious about what they write it should not be anonymous.
- Is it dated? Documents tend to evolve as the technology advances. Be sure you are reading the latest version. Internet search engines can help you here.
- Are there any followups? Be sure to check any followups or comments to what you read, otherwise you might miss a warning or a correction.

If you keep this in mind you should not fall for too many of the scams that circulate on the net, from get well–cards for Craig Shergold, chain letters to the more recent problems of Trojans for Linux that tries to trick you into mailing off your password.

There is a number of FAQs available that deals with more serious research method topics and you can also see a comprehensive on–line version.

Finally, do not forget the Linux Documentation Project site that coordinates documentation for Linux. Updates and new documents are issued regularly, reflecting the development in the field.

Quoted from someone's signature:

Be alert! The world needs more lerts.