Software Release Practice HOWTO

Eric Steven Raymond

Thyrsus Enterprises

esr@thyrsus.com

Copyright © 2000 by Eric S. Raymond

Revision History

Revision 3.4	2002−01−04	Revised by: esr
More about good patching practice.		
Revision 3.3	2001−08−16	Revised by: esr
New section about how to send good patches.		
Revision 3.2	2001−07−11	Revised by: esr
Note about not relying on proprietary components.		
Revision 3.1	2001−02−22	Revised by: esr
LDP Styleguide fixes.		
Revision 3.0	2000−08−12	Revised by: esr
First DocBook version. Advice on SourceForge and a major section on documentation practice added.		

This HOWTO describes good release practices for Linux and other open−source projects. By following these practices, you will make it as easy as possible for users to build your code and use it, and for other developers to understand your code and cooperate with you to improve it.

This document is a must−read for novice developers. Experienced developers should review it when they are about to release a new project. It will be revised periodically to reflect the evolution of good−practice standards.

Table of Contents

1. Introduction....1
1.1. Why this document?....1
1.2. New versions of this document....1

2. Good patching practice....2
2.1. Do send patches, don't send whole archives or files....2
2.2. Send patches against the current version of the code....2
2.3. Don't include patches for generated files....2
2.4. Don't send patch bands that just tweak RCS or SCCS $−symbols....3
2.5. Do use −c or −u format, don't use the default (−e) format....3
2.6. Do include documentation with your patch....3
2.7. Do include an explanation with your patch....4
2.8. Do include useful comments in your code....4

3. Good project− and archive− naming practice....6
3.1. Use GNU−style names with a stem and major.minor.patch numbering....6
3.2. But respect local conventions where appropriate....7
3.3. Try hard to choose a name prefix that is unique and easy to type....8

4. Good licensing and copyright practice: the theory....9
4.1. Open source and copyrights....9
4.2. What qualifies as open source....9

5. Good licensing and copyright practice: the practice....11
5.1. Make yourself or the FSF the copyright holder....11
5.2. Use a license conformant to the Open Source Definition....11
5.3. Don't write your own license if you can possibly avoid it....11

6. Good development practice....12
6.1. Write either pure ANSI C or a portable scripting language....12
6.2. Don't rely on proprietary code....12
6.3. Follow good C portability practices....12
6.4. Use autoconf/automake/autoheader....12
6.5. Sanity−check your code before release....13
6.6. Sanity−check your documentation and READMEs before release....13

7. Good distribution−making practice....14
7.1. Make sure tarballs always unpack into a single new directory....14
7.2. Have a README....14
7.3. Respect and follow standard file naming practices....14
7.4. Design for Upgradability....16
7.5. Provide RPMs....16

8. Good documentation practice....17
8.1. Good practice in the present....17
8.2. Good practice for the future....18

9. Good communication practice....20

Table of Contents

9.1. Announce to c.o.l.a and Freshmeat....20
9.2. Announce to a relevant topic newsgroup....20
9.3. Have a website....20
9.4. Host project mailing lists....20
9.5. Release to major archives....21

10. Good project−management practice....22

1. Introduction

1.1. Why this document?

There is a large body of good−practice traditions for open−source code that helps other people port, use, and cooperate with developing it. Some of these conventions are traditional in the Unix world and predate Linux; others have developed recently in response to particular new tools and technologies such as the World Wide Web.

This document will help you learn good practice. It is organized into topic sections, each containing a series of checklist items. Think of these as a pre−flight checklist for your patch or software distribution.

1.2. New versions of this document

This document will be posted monthly to the newsgroups comp.os.linux.answers. You can also view the latest version of this HOWTO on the World Wide Web via the URL http://www.linuxdoc.org/LDP/HOWTO/Software−Release−Practice.html.

Feel free to mail any questions or comments about this HOWTO to Eric S. Raymond, <esr@snark.thyrsus.com>.

---

2. Good patching practice

Most people get involved in open−source software by writing patches for other peoples' software before releasing projects of their own. Suppose you've written a set of source−code changes for someone else's baseline code. Now put yourself in that person's shoes. How is he to judge whether to include the patch?

The truth is that it is very difficult to judge the quality of code. So developers tend to evaluate patches by the quality of the submission. They look for clues in the submitter's style and communications behavior instead indications that the person has been in their shoes and understands what it's like to have to evaluate and merge an incoming patch.

This is actually a pretty reliable proxy for code quality. In many years of dealing with patches from many hundreds of strangers, I have only seldom seen a patch that was thoughtfully presented and respectful of my time but technically bogus. On the other hand, experience teaches me that patches which look careless or are packaged in a lazy and inconsiderate way are very likely to actually *be* bogus.

Here are some tips on how to get your patch accepted:

2.1. Do send patches, don't send whole archives or files

If your change includes a new file that doesn't exist in the code, then of course you have to send the whole file. But if you're modifying an already−existing file, don't send the whole file. Send a diff instead; specifically, send the output of the diff(1) command run to compare the baseline distributed version against your modified version.

The diff command and its dual, patch(1) (which automatically applies a diff to a baseline file) are the most basic tools of open−source development. Diffs are better than whole files because the developer you're sending a patch to may have changed the baseline version since you got your copy. By sending him a diff you save him the effort of separating your changes from his; you show respect for his time.

2.2. Send patches against the current version of the code.

It is both counterproductive and rude to send a maintainer patches against the code as it existed several releases ago, and expect him to do all the work of determining which changes duplicate things he have since done, versus which things are actually novel in your patch.

As a patch submitter, it is *your* responsibility to track the state of the source and send the maintainer a minimal patch that expresses what you want done to the main−line codebase. That means sending a patch against the current version.

2.3. Don't include patches for generated files.

Before you send your patch, walk through it and delete any patch bands for files in it that are going to be automatically regenerated once he applies the patch and remakes. The classic examples of this error are C files generated by Bison or Flex.

These days the most common mistake of this kind is sending a diff with a huge band that is nothing but changebars between your **configure** script and his. This file is generated by **autoconf**.

This is inconsiderate. It means your recipient is put to the trouble of separating the real content of the patch from a lot of bulky noise. It's a minor error, not as important as some of the things we'll get to further on but it will count against you.

2.4. Don't send patch bands that just tweak RCS or SCCS $−symbols.

Some people put special tokens in their source files that are expanded by the version−control system when the file is checked in: the $Id: Software−Release−Practice−HOWTO.sgml,v 1.15 2002/02/04 19:24:34 esr Exp $ construct used by RCS and CVS, for example.

If you're using a local version−control system yourself, your changes may alter these tokens. This isn't really harmful, because when your recipient checks his code back in after applying your patch they'll get re−expanded based on *his* version−control status. But those extra patch bands are noise. They're distracting. It's more considerate not to send them.

This is another minor error. You'll be forgiven for it if you get the big things right. But you want to avoid it anyway.

2.5. Do use −c or −u format, don't use the default (−e) format

The default (−e) format of diff(1) is very brittle. It doesn't include any context, so the patch tool can't cope if any lines have been inserted or deleted in the baseline code since you took the copy you modified.

Getting an −e diff is annoying, and suggests that the sender is either an extreme newbie, careless, or clueless. Most such patches get tossed out without a second thought.

2.6. Do include documentation with your patch

This is very important. If your patch makes a user−visible addition or change to the software's features, *include changes to the appropriate man pages and other documentation files in your patch*. Do *not* assume that the recipient will be happy to document your code for you, or else to have undocumented features lurking in the code.

Documenting your changes well demonstrates some good things. First, it's considerate to the person you are trying to persuade. Second, it shows that you understand the ramifications of your change well enough to explain it to somebody who can't see the code. Third, it demonstrates that you care about the people who will ultimately use the software.

Good documentation is usually the most visible sign of what separates a solid contribution from a quick and dirty hack. If you take the time and care necessary to produce it, you'll find you're already 85% of the way to having your patch accepted with most developers.

2.7. Do include an explanation with your patch

Your patch should include cover notes explaining why you think the patch is necessary or useful. This is explanation directed not to the users of the software but to the maintainer to whom you are sending the patch.

The note can be short in fact, some of the most effective cover notes I've ever seen just said "See the documentation updates in this patch". But it should show the right attitude.

The right attitude is helpful, respectful of the maintainer's time, quietly confident but unassuming. It's good to display understanding of the code you're patching. It's good to show that you can identify with the maintainer's problems. It's also good to be up front about any risks you perceive in applying the patch. Here are some examples of the sorts of explanatory comments that I like to see in cover notes:

"I've seen two problems with this code, X and Y. I fixed problem X, but I didn't try addressing problem Y because I don't think I understand the part of the code that I believe is involved."

"Fixed a core dump that can happen when one of the foo inputs is too long. While I was at it, I went looking for similar overflows elsewhere. I found a possible one in blarg.c, near line 666. Are you sure the sender can't generate more than 80 characters per transmission?"

"Have you considered using the Foonly algorithm for this problem? There is a good implementation at <http://www.somesite.com/~jsmith/foonly.html>."

"This patch solves the immediate problem, but I realize it complicates the memory allocation in an unpleasant way. Works for me, but you should probably test it under heavy load before shipping. "

"This may be featurititis, but I'm sending it anyway. Maybe you'll know a cleaner way to implement the feature."

2.8. Do include useful comments in your code

Usually as a maintainer, I will want to have strong confidence that I understand your changes before merging them in. This isn't an invariable rule; if you have a track record of good work, with me I may just run a casual eyeball over the changes before checking them in semi−automatically. But everything you can do to help me understand your code and decrease my uncertainty increases your chances that I will accept your patch.

Good comments in your code help me understand it. Bad comments don't.

Here's an example of a bad comment:

```
/* norman newbie fixed this 13 Aug 2001 */
```

This conveys no information. It's nothing but a muddy territorial bootprint you're planting in the middle of my code. If I take your patch (which you've made less likely) I'll almost certainly strip out this comment. If you want a credit, include a patch band for the project `NEWS` or `HISTORY` file. I'll probably take that.

Here's an example of a good comment:

```
/*
 * This conditional needs to be guarded so that crunch_data()
 * never gets passed a NULL pointer.  <norman_newbie@foosite.com>
 */
```

This comment shows that you understand not only my code but the kind of information that I need to have confidence in your changes. This kind of comment *gives* me confidence in your changes.

---

3. Good project– and archive– naming practice

As the load on maintainers of archives like Metalab, the PSA site and CPAN increases, there is an increasing trend for submissions to be processed partly or wholly by programs (rather than entirely by a human).

This makes it more important for project and archive–file names to fit regular patterns that computer programs can parse and understand.

3.1. Use GNU–style names with a stem and major.minor.patch numbering.

It's helpful to everybody if your archive files all have GNU–like names –– all–lower–case alphanumeric stem prefix, followed by a dash, followed by a version number, extension, and other suffixes.

Let's suppose you have a project you call `foobar' at version 1, release 2, level 3. If it's got just one archive part (presumably the sources), here's what its names should look:

foobar–1.2.3.tar.gz

> The source archive

foobar.lsm

> The LSM file (assuming you're submitting to Metalab).

Please *don't* use these:

foobar123.tar.gz

> This looks to many programs like an archive for a project called`foobar123' with no version number.

foobar1.2.3.tar.gz

> This looks to many programs like an archive for a project called `foobar1' at version 2.3.

foobar–v1.2.3.tar.gz

> Many programs think this goes with a project called `foobar–v1'.

foo_bar–1.2.3.tar.gz

> The underscore is hard for people to speak, type, and remember.

FooBar–1.2.3.tar.gz

> Unless you *like* looking like a marketing weenie. This is also hard for people to speak, type, and remember.

If you have to differentiate between source and binary archives, or between different kinds of binary, or express some kind of build option in the file name, please treat that as a file extension to go *after* the version number. That is, please do this:

foobar−1.2.3.src.tar.gz

: sources

foobar−1.2.3.bin.tar.gz

: binaries, type not specified

foobar−1.2.3.bin.ELF.tar.gz

: ELF binaries

foobar−1.2.3.bin.ELF.static.tar.gz

: ELF binaries statically linked

foobar−1.2.3.bin.SPARC.tar.gz

: SPARC binaries

Please *don't* use names like `foobar−ELF−1.2.3.tar.gz', because programs have a hard time telling type infixes (like `−ELF') from the stem.

A good general form of name has these parts in order:

1. project prefix
2. dash
3. version number
4. dot
5. "src" or "bin" (optional)
6. dot or dash (dot preferred)
7. binary type and options (optional)
8. archiving and compression extensions

3.2. But respect local conventions where appropriate

Some projects and communities have well−defined conventions for names and version numbers that aren't necessarily compatible with the above advice. For instance, Apache modules are generally named like mod_foo, and have both their own version number and the version of Apache with which they work. Likewise, Perl modules have version numbers that can be treated as floating point numbers (e.g., you might see 1.303 rather than 1.3.3), and the distributions are generally named Foo−Bar−1.303.tar.gz for version 1.303 of module Foo::Bar. (Perl itself, on the other hand, switched to using the conventions described in this document in late 1999.)

Look for and respect the conventions of specialized communities and developers; for general use, follow the above guidelines.

3.3. Try hard to choose a name prefix that is unique and easy to type

The stem prefix should be common to all a project's files, and it should be easy to read, type, and remember. So please don't use underscores. And don't capitalize or BiCapitalize without extremely good reason −− it messes up the natural human−eyeball search order and looks like some marketing weenie trying to be clever.

It confuses people when two different projects have the same stem name. So try to check for collisions before your first release. Two good places to check are the index file of Metalab and the appindex at Freshmeat. Another good place to check is SourceForge; do a name search there.

4. Good licensing and copyright practice: the theory

The license you choose defines the social contract you wish to set up among your co−developers and users. The copyright you put on the software will function mainly as a legal assertion of your right to set license terms on the software and derivative works of the software.

4.1. Open source and copyrights

Anything that is not public domain has a copyright, possibly more than one. Under the Berne Convention (which has been U.S. law since 1978), the copyright does not have to be explicit. That is, the authors of a work hold copyright even if there is no copyright notice.

Who counts as an author can be very complicated, especially for software that has been worked on by many hands. This is why licenses are important. By setting out the terms under which material can be used, they grant rights to the users that protect them from arbitrary actions by the copyright holders.

In proprietary software, the license terms are designed to protect the copyright. They're a way of granting a few rights to users while reserving as much legal territory is possible for the owner (the copyright holder). The copyright holder is very important, and the license logic so restrictive that the exact technicalities of the license terms are usually unimportant.

In open−source software, the situation is usually the exact opposite; the copyright exists to protect the license. The only rights the copyright holder always keeps are to enforce the license. Otherwise, only a few rights are reserved and most choices pass to the user. In particular, the copyright holder cannot change the terms on a copy you already have. Therefore, in open−source software the copyright holder is almost irrelevant −− but the license terms are very important.

Normally the copyright holder of a project is the current project leader or sponsoring organization. Transfer of the project to a new leader is often signaled by changing the copyright holder. However, this is not a hard and fast rule; many open−source projects have multiple copyright holders, and there is no instance on record of this leading to legal problems.

Some projects choose to assign copyright to the Free Software Foundation, on the theory that it has an interest in defending open source and lawyers available to do it.

4.2. What qualifies as open source

For licensing purposes, we can distinguish several different kinds of rights that a license may convey. Rights to *copy and redistribute*, rights to *use*, rights to *modify for personal use*, and rights to *redistribute modified copies*. A license may restrict or attach conditions to any of these rights.

The Open Source Initiative is the result of a great deal of thought about what makes software ``open source'' or (in older terminology) ``free''. Its constraints on licensing require that:

1. An unlimited right to copy be granted.

2. An unlimited right to use be granted.
3. An unlimited right to modify for personal use be granted.

The guidelines prohibit restrictions on redistribution of modified binaries; this meets the needs of software distributors, who need to be able to ship working code without encumbrance. It allows authors to require that modified sources be redistributed as pristine sources plus patches, thus establishing the author's intentions and an ``audit trail'' of any changes by others.

The OSD is the legal definition of the `OSI Certified Open Source' certification mark, and as good a definition of ``free software'' as anyone has ever come up with. All of the standard licenses (MIT, BSD, Artistic, and GPL/LGPL) meet it (though some, like GPL, have other restrictions which you should understand before choosing it).

Note that licenses which allow noncommercial use only do *not* qualify as open−source licenses, even if they are decorated with ``GPL'' or some other standard license. They discriminate against particular occupations, persons, and groups. They make life too complicated for CD−ROM distributors and others trying to spread open−source software commercially.

---

5. Good licensing and copyright practice: the practice

Here's how to translate the theory above into practice:

5.1. Make yourself or the FSF the copyright holder

In some cases, if you have a sponsoring organization behind you with lawyers, you might wish to give copyright to that organization.

5.2. Use a license conformant to the Open Source Definition

The Open Source Definition is the community gold standard for licenses. The OSD is not a license itself; rather, it defines a minimum set of rights that a license must guarantee in order to be considered an open−source license. The OSD, and supporting materials, may be found at the web site of the Open Source Initiative.

5.3. Don't write your own license if you can possibly avoid it.

The widely−known OSD−conformant licenses have well−established interpretive traditions. Developers (and, to the extent they care, users) know what they imply, and have a reasonable take on the risks and tradeoffs they involve. Therefore, use one of the standard licenses carried on the OSI site if at all possible.

If you must write your own license, be sure to have it certified by OSI. This will avoid a lot of argument and overhead. Unless you've been through it, you have no idea how nasty a licensing flamewar can get; people become passionate because the licenses are regarded as almost−sacred covenants touching the core values of the open−source community.

Furthermore, the presence of an established interpretive tradition may prove important if your license is ever tested in court. At time of writing (early 2002) there is no case law either supporting or invalidating any open−source license. However, it is a legal doctrine (at least in the U.S., and probably in other common−law countries such as England and the rest of the British Commonwealth) that courts are supposed to interpret licenses and contracts according to the expectations and practices of the community in which they originated.

6. Good development practice

Most of these are concerned with ensuring portability, not only across Linuxes but to other Unixes as well. Being portable to other Unixes is not just a worthy form of professionalism and hackerly politeness, it's valuable insurance against future changes in Linux itself.

Finally, other people *will* try to build your code on non−Linux systems; portability minimizes the number of annoying perplexed email messages you will get.

6.1. Write either pure ANSI C or a portable scripting language

For portability and stability, you should write either in ANSI C or a scripting language that is guaranteed portable because it has just one cross−platform implementation.

Scripting languages that qualify include Python, Perl, Tcl, Emacs Lisp, and PHP. Plain old shell does *not* qualify; there are too many different implementations with subtle idiosyncracies, and the shell environment is subject to disruption by user customizations such as shell aliases.

Java holds promise as a portable language, but the Linux−available implementations are still scratchy and poorly integrated with Linux. Java is still a bleeding−edge choice, though one likely to become more popular as it matures.

6.2. Don't rely on proprietary code

Don't rely on proprietary languages, libraries, or other code. In the open−source community this is considered rude. Open−source developers don't trust what they can't see the source code of.

6.3. Follow good C portability practices

If you are writing C, do feel free to use the full ANSI features −− including function prototypes, which will help you spot cross−module inconsistancies. The old−style K&R compilers are history.

On the other hand, do *not* assume that GCC−specific features such as the `−pipe' option or nested functions are available. These will come around and bite you the second somebody ports to a non−Linux, non−GCC system.

6.4. Use autoconf/automake/autoheader

If you're writing C, use autoconf/automake/autoheader to handle portability issues, do system−configuration probes, and tailor your makefiles. People building from sources today expect to be able to type "configure; make" and get a clean build −− and rightly so.

6.5. Sanity−check your code before release

If you're writing C, test−compile with −Wall and clean up the errors at least once before each release. This catches a surprising number of errors. For real thoroughness, compile with −pedantic as well.

For Python projects, the PyChecker program can be a useful check. It's not out of beta yet, but nevertheless often catches nontrivial errors.

If you're writing Perl, check your code with perl −c (and maybe −T, if applicable). Use perl −w and 'use strict' religiously. (See the Perl documentation for discussion.)

6.6. Sanity−check your documentation and READMEs before release

Run a spell−checker on them. If you look like you can't spell and don't care, people will assume your code is sloppy and careless too.

7. Good distribution–making practice

These guidelines describe how your distribution should look when someone downloads, retrieves and unpacks it.

7.1. Make sure tarballs always unpack into a single new directory

The single most annoying mistake newbie developers make is to build tarballs that unpack the files and directories in the distribution into the current directory, potentially stepping on files already located there. *Never do this!*

Instead, make sure your archive files all have a common directory part named after the project, so they will unpack into a single top–level directory directly *beneath* the current one.

Here's a makefile trick that, assuming your distribution directory is named `foobar' and SRC contains a list of your distribution files, accomplishes this.

```
foobar-$(VERS).tar.gz:
        @ls $(SRC) | sed s:^:foobar-$(VERS)/: >MANIFEST
        @(cd ..; ln -s foobar foobar-$(VERS))
        (cd ..; tar -czvf foobar/foobar-$(VERS).tar.gz `cat foobar/MANIFEST`)
        @(cd ..; rm foobar-$(VERS))
```

7.2. Have a README

Have a file called `README` or `READ.ME` that is a roadmap of your source distribution. By ancient convention, this is the first file intrepid explorers will read after unpacking the source.

Good things to have in the README include:

1. A brief description of the project.
2. A pointer to the project website (if it has one)
3. Notes on the developer's build environment and potential portability problems.
4. A roadmap describing important files and subdirectories.
5. Either build/installation instructions or a pointer to a file containing same (usually `INSTALL`).
6. Either a maintainers/credits list or a pointer to a file containing same (usually `CREDITS`).
7. Either recent project news or a pointer to a file containing same (usually `NEWS`).

7.3. Respect and follow standard file naming practices

Before even looking at the README, your intrepid explorer will have scanned the filenames in the top–level directory of your unpacked distribution. Those names can themselves convey information. By adhering to certain standard naming practices, you can give the explorer valuable clues about what to look in next.

Here are some standard top−level file names and what they mean. Not every distribution needs all of these.

README or READ.ME

the roadmap file, to be read first

INSTALL

configuration, build, and installation instructions

CREDITS

list of project contributers

NEWS

recent project news

HISTORY

project history

COPYING

project license terms (GNU convention)

LICENSE

project license terms

MANIFEST

list of files in the distribution

FAQ

plain−text Frequently−Asked−Questions document for the project

TAGS

generated tag file for use by Emacs or vi

Note the overall convention that filenames with all−caps names are human−readable metainformation about the package, rather than build components (TAGS is an exception to the first, but not to the second).

Having a FAQ can save you a lot of grief. When a question about the project comes up often, put it in the FAQ; then direct users to read the FAQ before sending questions or bug reports. A well−nurtured FAQ can decrease the support burden on the project maintainers by an order of magnitude or more.

Having a HISTORY or NEWS file with timestamps in it for each release is valuable. Among other things, it may help establish prior art if you are ever hit with a patent−infringement lawsuit (this hasn't happened to

anyone yet, but best to be prepared).

7.4. Design for Upgradability

Your software will change over time as you put out new releases. Some of these changes will not be backward−compatible. Accordingly, you should give serious thought to designing your installation layouts so that multiple installed versions of your code can coexist on the same system. This is especially important for libraries −− you can't count on all your client programs to upgrade in lockstep with your API changes.

The Emacs, Python, and Qt projects have a good convention for handling this; version−numbered directories. Here's how an installed Qt library hierarchy looks (${ver} is the version number):

```
/usr/lib/qt
/usr/lib/qt-${ver}
/usr/lib/qt-${ver}/bin          # Where you find moc
/usr/lib/qt-${ver}/lib          # Where you find .so
/usr/lib/qt-${ver}/include      # Where you find header files
```

With this organization, you can have multiple versions coexisting. Client programs have to specify the library version they want, but that's a small price to pay for not having the interfaces break on them.

7.5. Provide RPMs

The de−facto standard format for installable binary packages is that used by the Red Hat Package manager, RPM. It's featured in the most popular Linux distribution, and supported by effectively all other Linux distributions (except Debian and Slackware; and Debian can install from RPMs).

Accordingly, it's a good idea for your project site to provide installable RPMs as well as source tarballs.

It's also a good idea for you to include in your source tarball the RPM spec file, with a production that makes RPMs from it in your Makefile. The spec file should have the extension `.spec'; that's how the rpm −t option finds it in a tarball.

For extra style points, generate your spec file with a shellscript that automatically plugs in the correct version number by analyzing the Makefile or a version.h.

Note: if you supply source RPMs, use BuildRoot to make the program be built in /tmp or /var/tmp. If you don't, during the course of running the make install part of your build, the install will install the files in the real final places. This will happen even if there are file collisions, and even if you didn't want to install the package at all. When you're done, the files will have been installed and your system's RPM database will not know about it. Such badly behaved SRPMs are a minefield and should be eschewed.

8. Good documentation practice

The most important good documentation practice is to actually write some! Too many programmers omit this. But here are two good reasons to do it:

1. *Your documentation can be your design document.* The best time to write it is before you type a single line of code, while you're thinking out what you want to do. You'll find that the process of describing the way you want your program to work in natural language focuses your mind on the high−level questions about what it should do and how it should work. This may save you a lot of effort later.
2. *Your documentation is an advertisement for the quality of your code.* Many people take poor, scanty, or illiterate documentation for a program as a sign that the programmer is sloppy or careless of potential users' needs. Good documentation, on the other hand, conveys a message of intelligence and professionalism. If your program has to compete with other programs, better make sure your documentation is at least as good as theirs lest potential users write you off without a second look.

This HOWTO wouldn't be the place for a course on technical writing even if that were practical. So we'll focus here on the formats and tools available for composing and rendering documentation.

Though Unix and the open−source community have a long tradition of hosting powerful document−formatting tools, the plethora of different formats has meant that documentation has tended to be fragmented and difficult for users to browse or index in a coherent way. We'll summarize the uses, strengths, and weaknesses of the common documentation formats. Then we'll make some recommendations for good practice.

8.1. Good practice in the present

Here are the documentation markup formats now in widespread use among open−source developers. When we speak of "presentation" markup, we mean markup that controls the document's appearance explicitly (such as a font change). When we speak of "structural" markup, we mean markup that describes the logical structure of the document (like a section break or emphasis tag.) And when we speak of "indexing", we mean the process of extracting from a collection of documents a searchable collection of topic pointers that users can employ to reliably find material of interest across the entire collection.

man pages

The most most common format, inherited from Unix, a primitive form of presentation markup. man(1) command provides a pager and a stone−age search facility. No support for images or hyperlinks or indexing. Renders to Postscript for printing fairly well. Doesn't render to HTML at all well (essentially as flat text). Tools are preinstalled on all Linux systems.

Man page format is not bad for command summaries or short reference documents intended to jog the memory of an experienced user. It starts to creak under the strain for programs with complex interfaces and many options, and collapses entirely if you need to maintain a set of documents with rich cross−references (the markup has only weak and normally unused support for hyperlinks).

HTML

Increasingly common since the Web exploded in 1993−1994. Markup is partly structural, mostly presentation. Browseable through any web browser. Good support for images and hyperlinks. Limited built−in facilities for indexing, but good indexing and search−engine technologies exist and are widely deployed. Renders to Postscript for printing pretty well. HTML tools are now universally available.

HTML is very flexible and suitable for many kinds of documentation. Actually, it's *too* flexible; it shares with man page format the problem that it's hard to index automatically because a lot of the markup describes presentation rather than document structure.

Texinfo

Texinfo is the documentation format used by the Free Software Foundation. It's a set of macros on top of the powerful TeX formatting engine. Mostly structural, partly presentation. Browseable through Emacs or a standalone **info** program. Good support for hyperlinks, none for images. Good indexing for both print and on−line forms; when you install a Texinfo document, a pointer to it is automatically added to a browsable "dir" document listing all the Texinfo documents on your system. Renders to excellent Postscript and useable HTML. Texinfo tools are preinstalled on most Linux systems, and available at the Free Software Foundation website.

Texinfo is a good design, quite usable for typesetting books as well as small on−line documents, but like HTML it's a sort of amphibian −− the markup is part structural, part presentation, and the presentation part creates problems for rendering.

DocBook

DocBook is a large, elaborate markup format based on SGML (more recent versions on XML). Unlike the other formats described here it is entirely structural with no presentation markup. Excellent support for images and hyperlinks. Good support for indexing. Renders well to HTML, acceptably to Postscript for printing (quality is improving as the tools evolve). Tools and documentation are available at the DocBook website.

DocBook is excellent for large, complex documents; it was designed specifically to support technical manuals and rendering them in multiple output formats. Its drawbacks are complexity, a not entirely mature (though rapidly improving) toolset, and introductory−level documentation that is scanty and (too often) highly obfuscated.

8.2. Good practice for the future

In July of 2000 representatives from several important open−source project groups (including GNOME, KDE, the Free Software Foundation, the Linux Documentation Project, and the Open Source Initiative) held a summit conference in Monterey, California. The goal was to try and settle on common practices and common documentation interchange formats, so that a much richer and more unified body of documentation can evolve.

Concretely, the goal everyone has in view is to support a kind of documentation package which, when installed on a system, is immediately integrated into a rich system−wide index of documents in such a way that they can all be browsed through a uniform interface and searched as a unit. From the steps GNOME and KDE have already taken in this direction, it was already understood that this would require a structural rather

than presentation markup standard.

The meeting endorsed a trend which has been clear for a while; key open−source projects are moving or have already moved to DocBook as a master format for their documentation.

The participants also settled on using the `Dublin core' metadata format (an international standard developed by librarians concerned with the indexing of digital material) to support document indexing; details of that are still being worked out, and will probably result in some additions to the DocBook markup to support embedding Dublin Core metadata in DocBook documents.

The direction is clear; more use of DocBook, with auxiliary standards that support automatically indexing Docbook documents based on their index tags and Dublin core metadata. There are pieces still missing from this picture, but they will be filled in. The older presentation−based markups' days are numbered. (This HOWTO was moved to DocBook in August 2000.)

Thus, people starting new open−source projects will be ahead of the curve, and probably saving themselves a nasty conversion process later, if they go with DocBook as a master format from the beginning.

---

9. Good communication practice

Your software and documentation won't do the world much good if nobody but you knows it exists. Also, developing a visible presence for the project on the Internet will assist you in recruiting users and co−developers. Here are the standard ways to do that.

9.1. Announce to c.o.l.a and Freshmeat

Announce new releases to comp.os.linux.announce. Besides being widely read itself, this group is a major feeder for web−based what's−new sites like Freshmeat.

9.2. Announce to a relevant topic newsgroup

Find USENET topics group directly relevant to your application, and announce there as well. Post only where the *function* of the code is relevant, and exercise restraint.

If (for example) you are releasing a program written in Perl that queries IMAP servers, you should certainly post to comp.mail.imap. But you should probably not post to comp.lang.perl unless the program is also an instructive example of cutting−edge Perl techniques.

Your announcement should include the URL of a project website.

9.3. Have a website

If you intend try to build any substantial user or developer community around your project, it should have a website. Standard things to have on the website include:

- The project charter (why it exists, who the audience is, etc).
- Download links for the project sources.
- Instructions on how to join the project mailing list(s).
- A FAQ (Frequently Asked Questions) list.
- HTMLized versions of the project documentation
- Links to related and/or competing projects.

Some project sites even have URLs for anonymous access to the master source tree.

9.4. Host project mailing lists

It's standard practice to have a private development list through which project collaborators can communicate and exchange patches. You may also want to have an announcements list for people who want to be kept informed of the project's process.

If you are running a project named `foo'. your developer list might be foo−dev or foo−friends; your announcement list might be foo−announce.

9.5. Release to major archives

For the last several years, the Metalab archive has been the most important interchange location for Linux software.

Since it was launched in fall 1999, SourceForge has exploded in popularity. It is not just an archive and distribution site, though you can use it that way. It is an entire free project−hosting service that tries to offer a complete set of tools for open−source development groups −− web and archive space, mailing lists, bug−tracking, chat forums, CVS repositories, and other services.

Other important locations include:

- the Python Software Activity site (for software written in Python).
- the CPAN, the Comprehensive Perl Archive Network, (for software written in Perl).

10. Good project–management practice

Managing a project well when all the participants are volunteers presents some unique challenges. This is too large a topic to cover in a HOWTO. Fortunately, there are some useful white papers available that will help you understand the major issues.

For discussion of basic development organization and the release–early–release–often `bazaar mode', see The Cathedral and the Bazaar.

For discussion of motivational psychology, community customs, and conflict resolution, see Homesteading the Noosphere.

For discussion of economics and appropriate business models, see The Magic Cauldron.

These papers are not the last word on open–source development. But they were the first serious analyses to be written, and have yet to be superseded (though the author hopes they will be someday).