

# Committer's Guide

## The FreeBSD Documentation Project

Revision: [46036](#)

Copyright © 1999, 2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011, 2012, 2013, 2014 The FreeBSD Documentation Project

FreeBSD is a registered trademark of the FreeBSD Foundation.

Coverity is a registered trademark; Coverity Extend, Coverity Prevent and Coverity Prevent SQS are trademarks of Coverity, Inc.

IBM, AIX, OS/2, PowerPC, PS/2, S/390, and ThinkPad are trademarks of International Business Machines Corporation in the United States, other countries, or both.

Intel, Celeron, Centrino, Core, EtherExpress, i386, i486, Itanium, Pentium, and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

SPARC, SPARC64, and UltraSPARC are trademarks of SPARC International, Inc in the United States and other countries. SPARC International, Inc owns all of the SPARC trademarks and under licensing agreements allows the proper use of these trademarks by its members.

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this document, and the FreeBSD Project was aware of the trademark claim, the designations have been followed by the “™” or the “®” symbol.

2014-11-29 by rodrigc.

## Abstract

This document provides information for the FreeBSD committer community. All new committers should read this document before they start, and existing committers are strongly encouraged to review it from time to time.

Almost all FreeBSD developers have commit rights to one or more repositories. However, a few developers do not, and some of the information here applies to them as well. (For instance, some people only have rights to work with the Problem Report database). Please see [Section 19, “Issues Specific to Developers Who Are Not Committers”](#) for more information.

This document may also be of interest to members of the FreeBSD community who want to learn more about how the project works.

## Table of Contents

1. Administrative Details .....	2
2. OpenPGP Keys for FreeBSD .....	3
3. Kerberos and LDAP web Password for FreeBSD Cluster .....	5
4. Commit Bit Types .....	6
5. Subversion Primer .....	7
6. Setup, Conventions, and Traditions .....	29
7. Commit Log Messages .....	33
8. Preferred License for New Files .....	36
9. Keeping Track of Licenses Granted to the FreeBSD Project .....	38
10. Developer Relations .....	39
11. If in Doubt... ..	40
12. Bugzilla .....	40
13. Who's Who .....	41
14. SSH Quick-Start Guide .....	42
15. Coverity® Availability for FreeBSD Committers .....	43
16. The FreeBSD Committers' Big List of Rules .....	43
17. Support for Multiple Architectures .....	50
18. Ports Specific FAQ .....	53
19. Issues Specific to Developers Who Are Not Committers .....	61
20. Information About Google Analytics .....	61
21. Perks of the Job .....	62
22. Miscellaneous Questions .....	63

## 1. Administrative Details

<i>Login Methods</i>	<a href="#">ssh(1)</a> , protocol 2 only
<i>Main Shell Host</i>	freefall.FreeBSD.org
<i>src/ Subversion Root</i>	svn+ssh:// svn.FreeBSD.org /base (see also <a href="#">Section 5.2.3, “RELENG_* Branches and General Layout”</a> ).
<i>doc/ Subversion Root</i>	svn+ssh:// svn.FreeBSD.org /doc (see also <a href="#">Section 5.2.4, “FreeBSD Documentation Project Branches and Layout”</a> ).
<i>ports/ Subversion Root</i>	svn+ssh:// svn.FreeBSD.org /ports (see also <a href="#">Section 5.2.5, “FreeBSD Ports Tree Branches and Layout”</a> ).
<i>Internal Mailing Lists</i>	developers (technically called all-developers), doc-developers, doc-committers, ports-developers, ports-committers, src-developers, src-committers. (Each project repository has its own -developers

	and -committers mailing lists. Archives for these lists may be found in files <code>/home/mail/ repository-name-developers-archive</code> and <code>/home/mail/ repository-name-committers-archive</code> on the FreeBSD.org cluster.)
Core Team monthly reports	<code>/home/core/public/monthly-reports</code> on the FreeBSD.org cluster.
Ports Management Team monthly reports	<code>/home/portmgr/public/monthly-reports</code> on the FreeBSD.org cluster.
Noteworthy src/SVN Branches	<code>stable/8</code> (8.X-STABLE), <code>stable/9</code> (9.X-STABLE), <code>stable/10</code> (10.X-STABLE), <code>head</code> (-CURRENT)

`ssh(1)` is required to connect to the project hosts. For more information, see [Section 14, “SSH Quick-Start Guide”](#).

Useful links:

- [FreeBSD Project Internal Pages](#)
- [FreeBSD Project Hosts](#)
- [FreeBSD Project Administrative Groups](#)

## 2. OpenPGP Keys for FreeBSD

Cryptographic keys conforming to the OpenPGP (*Pretty Good Privacy*) standard are used by the FreeBSD project to authenticate committers. Messages carrying important information like public SSH keys can be signed with the OpenPGP key to prove that they are really from the committer. See [PGP & GPG: Email for the Practical Paranoid by Michael Lucas](#) and [http://en.wikipedia.org/wiki/Pretty\\_Good\\_Privacy](http://en.wikipedia.org/wiki/Pretty_Good_Privacy) for more information.

### 2.1. Creating a Key

Existing keys can be used, but should be checked with `doc/head/share/pgp-keys/checkkey.sh` first.

For those who do not yet have an OpenPGP key, or need a new key to meet FreeBSD security requirements, here we show how to generate one.

1. Install `security/gnupg`. Enter these lines in `~/.gnupg/gpg.conf` to set minimum acceptable defaults:

```
fixed-list-mode
keyid-format 0xlong
personal-digest-preferences SHA512 SHA384 SHA256 SHA224
```

```
default-preference-list SHA512 SHA384 SHA256 SHA224 AES256 ɔ
AES192 AES CAST5 BZIP2 ZLIB ZIP Uncompressed
use-agent
verify-options show-uid-validity
list-options show-uid-validity
sig-notation issuer-fpr@notations.openpgp.fifthhorseman.net=%g
cert-digest-algo SHA512
```

2. Generate a key:

```
% gpg --gen-key
gpg (GnuPG) 2.0.22; Copyright (C) 2013 Free Software Foundation, ɔ
Inc.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

Warning: using insecure memory!
Please select what kind of key you want:
  (1) RSA and RSA (default)
  (2) DSA and Elgamal
  (3) DSA (sign only)
  (4) RSA (sign only)
Your selection? 1
RSA keys may be between 1024 and 4096 bits long.
What keysize do you want? (2048) 2048 ❶
Requested keysize is 2048 bits
Please specify how long the key should be valid.
  0 = key does not expire
  <n> = key expires in n days
  <n>w = key expires in n weeks
  <n>m = key expires in n months
  <n>y = key expires in n years
Key is valid for? (0) 3y ❷
Key expires at Wed Nov  4 17:20:20 2015 MST
Is this correct? (y/N) y

GnuPG needs to construct a user ID to identify your key.

Real name: Chucky Daemon ❸
Email address: notreal@example.com
Comment:
You selected this USER-ID:
  "Chucky Daemon <notreal@example.com> "

Change (N)ame, (C)omment, (E)mail or (O)kay/(Q)uit? ɔ
You need a Passphrase to protect your secret key.
```

- ❶ 2048-bit keys with a three-year expiration provide adequate protection at present (2013-12). <http://danielpocock.com/rsa-key-sizes-2048-or-4096-bits> describes the situation in more detail.
- ❷ A three year key lifespan is short enough to obsolete keys weakened by advancing computer power, but long enough to reduce key management problems.

- ③ Use your real name here, preferably matching that shown on government-issued ID to make it easier for others to verify your identity. Text that may help others identify you can be entered in the `Comment` section.

After the email address is entered, a passphrase is requested. Methods of creating a secure passphrase are contentious. Rather than suggest a single way, here are some links to sites that describe various methods: <http://world.std.com/~reinhold/diceware.html>, <http://www.iusmentis.com/security/passphrasefaq/>, <http://xkcd.com/936/>, <http://en.wikipedia.org/wiki/Passphrase>.

Protect your private key and passphrase. If either the private key or passphrase may have been compromised or disclosed, immediately notify [<accounts@FreeBSD.org>](mailto:accounts@FreeBSD.org) and revoke the key.

Committing the new key is shown in [Procedure 1, “Steps for New Committers”](#).

### 3. Kerberos and LDAP web Password for FreeBSD Cluster

The FreeBSD cluster requires a Kerberos password to access certain services. The Kerberos password also serves as the LDAP web password, since LDAP is proxying to Kerberos in the cluster. Some of the services which require this include:

- [Bugzilla](#)
- [Jenkins](#)

To reset a Kerberos password in the FreeBSD cluster using a random password generator:

```
% ssh kpasswd.freebsd.org
```



#### Note

This must be done from a machine outside of the FreeBSD.org cluster.

A Kerberos password can also be set manually by logging into `freefall.FreeBSD.org` and running:

```
% kpasswd
```

## 4. Commit Bit Types

The FreeBSD repository has a number of components which, when combined, support the basic operating system source, documentation, third party application ports infrastructure, and various maintained utilities. When FreeBSD commit bits are allocated, the areas of the tree where the bit may be used are specified. Generally, the areas associated with a bit reflect who authorized the allocation of the commit bit. Additional areas of authority may be added at a later date: when this occurs, the committer should follow normal commit bit allocation procedures for that area of the tree, seeking approval from the appropriate entity and possibly getting a mentor for that area for some period of time.

<i>Committer Type</i>	<i>Responsible</i>	<i>Tree Components</i>
src	core@	src/, doc/ subject to appropriate review
doc	doceng@	doc/, src/ documentation
ports	portmgr@	ports/

Commit bits allocated prior to the development of the notion of areas of authority may be appropriate for use in many parts of the tree. However, common sense dictates that a committer who has not previously worked in an area of the tree seek review prior to committing, seek approval from the appropriate responsible party, and/or work with a mentor. Since the rules regarding code maintenance differ by area of the tree, this is as much for the benefit of the committer working in an area of less familiarity as it is for others working on the tree.

Committers are encouraged to seek review for their work as part of the normal development process, regardless of the area of the tree where the work is occurring.

### 4.1. Policy for doc/ Committer Activity in src/

- doc committers may commit documentation changes to src files, such as man pages, READMEs, fortune databases, calendar files, and comment fixes without approval from a src committer, subject to the normal care and tending of commits.
- doc committers may commit minor src changes and fixes, such as build fixes, small features, etc, with an "Approved by" from a src committer.
- doc committers may seek an upgrade to a src commit bit by acquiring a mentor, who will propose the doc committer to core. When approved, they will be added to 'access' and the normal mentoring period will ensue, which will involve a continuing of "Approved by" for some period.
- "Approved by" is only acceptable from non-mentored src committers -- mentored committers can provide a "Reviewed by" but not an "Approved by".

## 5. Subversion Primer

It is assumed that you are already familiar with the basic operation of Subversion. If not, start by reading the [Subversion Book](#).

### 5.1. Introduction

The FreeBSD source repository switched from CVS to Subversion on May 31st, 2008. The first real SVN commit is *r179447*.

The FreeBSD doc/www repository switched from CVS to Subversion on May 19th, 2012. The first real SVN commit is *r38821*.



#### Note

Part of the doc/www CVS to SVN conversion included an infrastructural change to the build process. The most notable change is the location of the FreeBSD website `www` tree, which has been moved from `www/lang/` to `head/lang/htdocs/`.

The FreeBSD ports repository switched from CVS to Subversion on July 14th, 2012. The first real SVN commit is *r300894*.

Subversion has a number of features that should make developers' lives easier. The most important advantage to Subversion (and the reason why FreeBSD switched) is that it handles branches and merging much better than CVS does. Some of the principal differences are:

Subversion can be installed from the FreeBSD Ports Collection by issuing these commands:

```
# pkg install subversion
```

### 5.2. Getting Started

There are a few ways to obtain a working copy of the tree from Subversion. This section will explain them.

#### 5.2.1. Direct Checkout

The first is to check out directly from the main repository. For the `src` tree, use:

```
% svn checkout svn+ssh://svn.freebsd.org/base/head /usr/src
```

For the doc tree, use:

```
% svn checkout svn+ssh://svn.freebsd.org/doc/head /usr/doc
```

For the ports tree, use:

```
% svn checkout svn+ssh://svn.freebsd.org/ports/head /usr/ports
```



### Note

Though the remaining examples in this document are written with the workflow of working with the `src` tree in mind, the underlying concepts are the same for working with the `doc` and the `ports` tree. Ports related Subversion operations are listed in [Section 18, “Ports Specific FAQ”](#).

The above command will check out a `CURRENT` source tree as `/usr/src/`, which can be any target directory on the local filesystem. Omitting the final argument of that command causes the working copy, in this case, to be named “head”, but that can be renamed safely.

`svn+ssh` means the SVN protocol tunnelled over SSH. The name of the server is `svn.freebsd.org`, `base` is the path to the repository, and `head` is the subdirectory within the repository.

If your FreeBSD login name is different from your login name on your local machine, you must either include it in the URL (for example `svn+ssh://jarjar@svn.freebsd.org/base/head`), or add an entry to your `~/.ssh/config` in the form:

```
Host svn.freebsd.org
  User jarjar
```

This is the simplest method, but it is hard to tell just yet how much load it will place on the repository. Subversion is much faster than CVS, however.



### Note

The `svn diff` does not require access to the server as SVN stores a reference copy of every file in the working copy. This, however, means that Subversion working copies are very large in size.



### 5.2.2. Checkout from a Mirror

Check out a working copy from a mirror by substituting the mirror's URL for `svn+ssh://svn.freebsd.org/base` . This can be an official mirror or a mirror maintained by using `svnsync`.

There is a serious disadvantage to this method: every time something is to be committed, a `svn relocate` to the master repository has to be done, remembering to `svn relocate` back to the mirror after the commit. Also, since `svn relocate` only works between repositories that have the same UUID, some hacking of the local repository's UUID has to occur before it is possible to start using it.

Unlike with CVS, the hassle of a local `svnsync` mirror probably is not worth it unless the network connectivity situation or other factors demand it. If it is needed, see the end of this chapter for information on how to set one up.

### 5.2.3. RELENG\_\* Branches and General Layout

In `svn+ssh://svn.freebsd.org/base` , *base* refers to the source tree. Similarly, *ports* refers to the ports tree, and so on. These are separate repositories with their own change number sequences, access controls and commit mail.

For the base repository, HEAD refers to the -CURRENT tree. For example, `head/bin/ls` is what would go into `/usr/src/bin/ls` in a release. Some key locations are:

- `/head/` which corresponds to HEAD, also known as -CURRENT.
- `/stable/n` which corresponds to RELENG\_*n*.
- `/releng/n.n` which corresponds to RELENG\_*n\_n*.
- `/release/n.n.n` which corresponds to RELENG\_*n\_n\_n*\_RELEASE .
- `/vendor*` is the vendor branch import work area. This directory itself does not contain branches, however its subdirectories do. This contrasts with the *stable*, *releng* and *release* directories.
- `/projects` and `/user` feature a branch work area, like in Perforce. As above, the `/user` directory does not contain branches itself.

### 5.2.4. FreeBSD Documentation Project Branches and Layout

In `svn+ssh://svn.freebsd.org/doc` , *doc* refers to the repository root of the source tree.

In general, most FreeBSD Documentation Project work will be done within the `head/` branch of the documentation source tree.

FreeBSD documentation is written and/or translated to various languages, each in a separate directory in the `head/` branch.

Each translation set contains several subdirectories for the various parts of the FreeBSD Documentation Project. A few noteworthy directories are:

- `/articles/` contains the source code for articles written by various FreeBSD contributors.
- `/books/` contains the source code for the different books, such as the FreeBSD Handbook.
- `/htdocs/` contains the source code for the FreeBSD website.

### 5.2.5. FreeBSD Ports Tree Branches and Layout

In `svn+ssh://svn.freebsd.org/ports`, `ports` refers to the repository root of the ports tree.

In general, most FreeBSD port work will be done within the `head/` branch of the ports tree which is the actual ports tree used to install software. Some other key locations are:

- `/branches/RELENG_n_n_n` which corresponds to `RELENG_n_n_n` is used to merge back security updates in preparation for a release.
- `/tags/RELEASE_n_n_n` which corresponds to `RELEASE_n_n_n` represents a release tag of the ports tree.
- `/tags/RELEASE_n_EOL` represents the end of life tag of a specific FreeBSD branch.

## 5.3. Daily Use

This section will explain how to perform common day-to-day operations with Subversion.

### 5.3.1. Help

SVN has built in help documentation. It can be accessed by typing the following command:

```
% svn help
```

Additional information can be found in the [Subversion Book](#).

### 5.3.2. Checkout

As seen earlier, to check out the FreeBSD head branch:

```
% svn checkout svn+ssh://svn.freebsd.org/base/head /usr/src
```

At some point, more than just `HEAD` will probably be useful, for instance when merging changes to `stable/7`. Therefore, it may be useful to have a partial checkout of the complete tree (a full checkout would be very painful).

To do this, first check out the root of the repository:

```
% svn checkout --depth=immediates svn+ssh://svn.freebsd.org/base
```

This will give base with all the files it contains (at the time of writing, just ROADMAP.txt ) and empty subdirectories for head, stable, vendor and so on.

Expanding the working copy is possible. Just change the depth of the various subdirectories:

```
% svn up --set-depth=infinity base/head
% svn up --set-depth=immediates base/release base/releeng base/stable
```

The above command will pull down a full copy of head, plus empty copies of every release tag, every releeng branch, and every stable branch.

If at a later date merging to 7-STABLE is required, expand the working copy:

```
% svn up --set-depth=infinity base/stable/7
```

Subtrees do not have to be expanded completely. For instance, expanding only stable/7/sys and then later expand the rest of stable/7:

```
% svn up --set-depth=infinity base/stable/7/sys
% svn up --set-depth=infinity base/stable/7
```

Updating the tree with svn update will only update what was previously asked for (in this case, head and stable/7; it will not pull down the whole tree.



### Note

Decreasing the depth of a working copy is not possible.

## 5.3.3. Anonymous Checkout

It is possible to anonymously check out the FreeBSD repository with Subversion. This will give access to a read-only tree that can be updated, but not committed back to the main repository. To do this, use the following command:

```
% svn co https://svn0.us-west.FreeBSD.org /base/head /usr/src
```

Select the closest mirror and verify the mirror server certificate from the list of [Subversion mirror sites](#).

## 5.3.4. Updating the Tree

To update a working copy to either the latest revision, or a specific revision:

```
% svn update
% svn update - r12345
```

### 5.3.5. Status

To view the local changes that have been made to the working copy:

```
% svn status
```

To show local changes and files that are out-of-date do:

```
% svn status --show-updates
```

### 5.3.6. Editing and Committing

Unlike Perforce, SVN does not need to be told in advance about file editing.

`svn commit` works like the equivalent CVS command. To commit all changes in the current directory and all subdirectories:

```
% svn commit
```

To commit all changes in, for example, `lib/libfetch/` and `usr/bin/fetch/` in a single operation:

```
% svn commit lib/libfetch usr/bin/fetch
```

There is also a commit wrapper for the ports tree to handle the properties and sanity checking your changes:

```
% /usr/ports/Tools/scripts/psvn commit
```

### 5.3.7. Adding and Removing Files



#### Note

Before adding files, get a copy of [auto-props.txt](#) (there is also a [ports tree specific version](#)) and add it to `~/.subversion/config` according to the instructions in the file. If you added something before reading this, use `svn rm --keep-local` for just added files, fix your config file and re-add them again. The initial config file is created when you first run a svn command, even something as simple as `svn help`.

Files are added to a SVN repository with `svn add`. To add a file named `foo`, edit it, then:

```
% svn add foo
```



## Note

Most new source files should include a `$FreeBSD$` string near the start of the file. On commit, svn will expand the `$FreeBSD$` string, adding the file path, revision number, date and time of commit, and the username of the committer. Files which cannot be modified may be committed without the `$FreeBSD$` string.

Files can be removed with `svn remove`:

```
% svn remove foo
```

Subversion does not require deleting the file before using `svn rm`, and indeed complains if that happens.

It is possible to add directories with `svn add`:

```
% mkdir bar  
% svn add bar
```

Although `svn mkdir` makes this easier by combining the creation of the directory and the adding of it:

```
% svn mkdir bar
```

Like files, directories are removed with `svn rm`. There is no separate command specifically for removing directories.

```
% svn rm bar
```

### 5.3.8. Copying and Moving Files

This command creates a copy of `foo.c` named `bar.c`, with the new file also under version control:

```
% svn copy foo.c bar.c
```

The example above is equivalent to:

```
% cp foo.c bar.c  
% svn add bar.c
```

To move and rename a file:

```
% svn move foo.c bar.c
```

### 5.3.9. Log and Annotate

`svn log` shows revisions and commit messages, most recent first, for files or directories. When used on a directory, all revisions that affected the directory and files within that directory are shown.

`svn annotate`, or equally `svn praise` or `svn blame`, shows the most recent revision number and who committed that revision for each line of a file.

### 5.3.10. Diffs

`svn diff` displays changes to the working copy. Diffs generated by SVN are unified and include new files by default in the diff output.

`svn diff` can show the changes between two revisions of the same file:

```
% svn diff -r179453:179454 ROADMAP.txt
```

It can also show all changes for a specific changeset. The following will show what changes were made to the current directory and all subdirectories in changeset 179454:

```
% svn diff -c179454 .
```

### 5.3.11. Reverting

Local changes (including additions and deletions) can be reverted using `svn revert`. It does not update out-of-date files, but just replaces them with pristine copies of the original version.

### 5.3.12. Conflicts

If an `svn update` resulted in a merge conflict, Subversion will remember which files have conflicts and refuse to commit any changes to those files until explicitly told that the conflicts have been resolved. The simple, not yet deprecated procedure is the following:

```
% svn resolved foo
```

However, the preferred procedure is:

```
% svn resolve --accept=working foo
```

The two examples are equivalent. Possible values for `--accept` are:

- `working`: use the version in your working directory (which one presumes has been edited to resolve the conflicts).

- **base**: use a pristine copy of the version you had before `svn update`, discarding your own changes, the conflicting changes, and possibly other intervening changes as well.
- **mine-full**: use what you had before `svn update`, including your own changes, but discarding the conflicting changes, and possibly other intervening changes as well.
- **theirs-full**: use the version that was retrieved when you did `svn update`, discarding your own changes.

## 5.4. Advanced Use

### 5.4.1. Sparse Checkouts

SVN allows *sparse*, or partial checkouts of a directory by adding `--depth` to a `svn checkout`.

Valid arguments to `--depth` are:

- **empty**: the directory itself without any of its contents.
- **files**: the directory and any files it contains.
- **immediates**: the directory and any files and directories it contains, but none of the subdirectories' contents.
- **infinity**: anything.

The `--depth` option applies to many other commands, including `svn commit`, `svn revert`, and `svn diff`.

Since `--depth` is sticky, there is a `--set-depth` option for `svn update` that will change the selected depth. Thus, given the working copy produced by the previous example:

```
% cd ~/freebsd
% svn update --set-depth=immediates .
```

The above command will populate the working copy in `~/freebsd` with `ROADMAP.txt` and empty subdirectories, and nothing will happen when `svn update` is executed on the subdirectories. However, the following command will set the depth for *head* (in this case) to infinity, and fully populate it:

```
% svn update --set-depth=infinity head
```

### 5.4.2. Direct Operation

Certain operations can be performed directly on the repository without touching the working copy. Specifically, this applies to any operation that does not require editing a file, including:

- log, diff
- mkdir
- remove, copy, rename
- propset, propedit, propdel
- merge

Branching is very fast. The following command would be used to branch RELENG\_8:

```
% svn copy svn+ssh://svn.freebsd.org/base/head svn+ssh://svn.freebsd.org/base/stable/8
```

This is equivalent to the following set of commands which take minutes and hours as opposed to seconds, depending on your network connection:

```
% svn checkout --depth=immediates svn+ssh://svn.freebsd.org/base
% cd base
% svn update --set-depth=infinity head
% svn copy head stable/8
% svn commit stable/8
```

### 5.4.3. Merging with SVN

This section deals with merging code from one branch to another (typically, from head to a stable branch).



#### Note

In all examples below, \$FSVN refers to the location of the FreeBSD Subversion repository, `svn+ssh://svn.freebsd.org/base/` .

#### 5.4.3.1. About Merge Tracking

From the user's perspective, merge tracking information (or mergeinfo) is stored in a property called `svn:mergeinfo` , which is a comma-separated list of revisions and ranges of revisions that have been merged. When set on a file, it applies only to that file. When set on a directory, it applies to that directory and its descendants (files and directories) except for those that have their own `svn:mergeinfo` .

It is *not* inherited. For instance, `stable/6/contrib/openpam/` does not implicitly inherit mergeinfo from `stable/6/`, or `stable/6/contrib/`. Doing so would make partial checkouts very hard to manage. Instead, mergeinfo is explicitly propagated down the tree. For merging something into `branch/foo/bar/` , the following rules apply:



1. If `branch/foo/bar/` does not already have a mergeinfo record, but a direct ancestor (for instance, `branch/foo/`) does, then that record will be propagated down to `branch/foo/bar/` before information about the current merge is recorded.
2. Information about the current merge will *not* be propagated back up that ancestor.
3. If a direct descendant of `branch/foo/bar/` (for instance, `branch/foo/bar/baz/`) already has a mergeinfo record, information about the current merge will be propagated down to it.

If you consider the case where a revision changes several separate parts of the tree (for example, `branch/foo/bar/` and `branch/foo/quux/`), but you only want to merge some of it (for example, `branch/foo/bar/`), you will see that these rules make sense. If mergeinfo was propagated up, it would seem like that revision had also been merged to `branch/foo/quux/`, when in fact it had not been.

#### 5.4.3.2. Selecting the Source and Target for `stable/10` and Newer

Starting with the `stable/10` branch, all merges should be merged to and committed from the root of the branch. All merges should look like:

```
% svn merge -c r123456 ^/head/ checkout
% svn commit checkout
```

Note that *checkout* should be a complete checkout of the branch to which the merge occurs.

Merges to `releng/` branches should always originate from the corresponding `stable/` branch. For example:

```
% svn merge -c r123456 ^/stable/10 releng/10.0
```

#### 5.4.3.3. Selecting the Source and Target for `stable/9` and Older

For `stable/9` and earlier, a different strategy was used, distributing mergeinfo around the tree so that merges could be performed without a complete checkout. This procedure proved extremely error-prone, with the convenience of partial checkouts for merges significantly outweighed by the complexity of picking mergeinfo targets. The below describes this now-obsolete procedure, which should be used *only for merges prior to `stable/10`*.

Because of mergeinfo propagation, it is important to choose the source and target for the merge carefully to minimise property changes on unrelated directories.

The rules for selecting the merge target (the directory that you will merge the changes to) can be summarized as follows:

1. Never merge directly to a file.

2. Never, ever merge directly to a file.
3. *Never, ever, ever* merge directly to a file.
4. Changes to kernel code should be merged to `sys/`. For instance, a change to the `ichwd(4)` driver should be merged to `sys/`, not `sys/dev/ichwd/`. Likewise, a change to the TCP/IP stack should be merged to `sys/`, not `sys/netinet/`.
5. Changes to code under `etc/` should be merged at `etc/`, not below it.
6. Changes to vendor code (code in `contrib/`, `crypto/` and so on) should be merged to the directory where vendor imports happen. For instance, a change to `crypto/openssl/util/` should be merged to `crypto/openssl/`. This is rarely an issue, however, since changes to vendor code are usually merged wholesale.
7. Changes to userland programs should as a general rule be merged to the directory that contains the Makefile for that program. For instance, a change to `usr.bin/xlint/arch/i386/` should be merged to `usr.bin/xlint/`.
8. Changes to userland libraries should as a general rule be merged to the directory that contains the Makefile for that library. For instance, a change to `lib/libc/gen/` should be merged to `lib/libc/`.
9. There may be cases where it makes sense to deviate from the rules for userland programs and libraries. For instance, everything under `lib/libpam/` is merged to `lib/libpam/`, even though the library itself and all of the modules each have their own Makefile.
10. Changes to manual pages should be merged to `share/man/man N/`, for the appropriate value of `N`.
11. Other changes to `share/` should be merged to the appropriate subdirectory and not to `share/` directly.
12. Changes to a top-level file in the source tree such as `UPDATING` or `Makefile.inc1` should be merged directly to that file rather than to the root of the whole tree. Yes, this is an exception to the first three rules.
13. When in doubt, ask.

If you need to merge changes to several places at once (for instance, changing a kernel interface and every userland program that uses it), merge each target separately, then commit them together. For instance, if you merge a revision that changed a kernel API and updated all the userland bits that used that API, you would merge the kernel change to `sys`, and the userland bits to the appropriate userland directories, then commit all of these in one go.

The source will almost invariably be the same as the target. For instance, you will always merge `stable/7/lib/libc/` from `head/lib/libc/`. The only exception would be when merging changes to code that has moved in the source branch but not in the parent branch. For instance, a change to `pkg(1)` would be merged from `bin/pkg/` in `head` to `usr.bin/pkg/` in `stable/7`.

#### 5.4.3.4. Preparing the Merge Target

Because of the mergeinfo propagation issues described earlier, it is very important that you never merge changes into a sparse working copy. You must always have a full checkout of the branch you will merge into. For instance, when merging from `HEAD` to `7`, you must have a full checkout of `stable/7`:

```
% cd stable/7
% svn up --set-depth=infinity
```

The target directory must also be up-to-date and must not contain any uncommitted changes or stray files.

#### 5.4.3.5. Identifying Revisions

Identifying revisions to be merged is a must. If the target already has complete mergeinfo, ask SVN for a list:

```
% cd stable/6/contrib/openpam
% svn mergeinfo --show-revs=eligible $FSVN/head/contrib/openpam
```

If the target does not have complete mergeinfo, check the log for the merge source.

#### 5.4.3.6. Merging

Now, let us start merging!

##### 5.4.3.6.1. The Principles

Say you would like to merge:

- revision `$R`
- in directory `$target` in stable branch `$B`
- from directory `$source` in `head`
- `$FSVN` is `svn+ssh://svn.freebsd.org/base`

Assuming that revisions `$P` and `$Q` have already been merged, and that the current directory is an up-to-date working copy of `stable/$B`, the existing mergeinfo looks like this:

```
% svn propget svn:mergeinfo -R $target
$target - /head/$source:$P,$Q
```

Merging is done like so:

```
% svn merge -cR $FSVN/head/$source $target
```

Checking the results of this is possible with `svn diff`.

The `svn:mergeinfo` now looks like:

```
% svn propget svn:mergeinfo -R $target
$target - head/$source:$P,$Q,$R
```

If the results are not exactly as shown, assistance may be required before committing as mistakes may have been made, or there may be something wrong with the existing `mergeinfo`, or there may be a bug in Subversion.

#### 5.4.3.6.2. Practical Example

As a practical example, consider the following scenario. The changes to `netmap.4` in `r238987` are to be merged from `CURRENT` to `9-STABLE`. The file resides in `head/share/man/man4`. According to [Section 5.4.3, “Merging with SVN”](#), this is also where to do the merge. Note that in this example all paths are relative to the top of the `svn` repository. For more information on the directory layout, see [Section 5.2.3, “RELENG\\_\\* Branches and General Layout”](#).

The first step is to inspect the existing `mergeinfo`.

```
% svn propget svn:mergeinfo -R stable/9/share/man/man4
```

Take a quick note of how it looks before moving on to the next step; doing the actual merge:

```
% svn merge -c r238987 svn+ssh://svn.freebsd.org/base/head/share/man/
man4 stable/9/share/man/man4
--- Merging r238987 into 'stable/9/share/man/man4':
U   stable/9/share/man/man4/netmap.4
--- Recording mergeinfo for merge of r238987 into
'stable/9/share/man/man4':
U   stable/9/share/man/man4
```

Check that the revision number of the merged revision has been added. Once this is verified, the only thing left is the actual commit.

```
% svn commit stable/9/share/man/man4
```

#### 5.4.3.6.3. Merging into the Kernel (sys/)

As stated above, merging into the kernel is different from merging in the rest of the tree. In many ways merging to the kernel is simpler because there is always the same merge target (`sys/`).

Once `svn merge` has been executed, `svn diff` has to be run on the directory to check the changes. This may show some unrelated property changes, but these can be ignored. Next, build and test the kernel, and, once the tests are complete, commit the code as normal, making sure that the commit message starts with “Merge `r226222` from head”, or similar.

### 5.4.3.7. Precautions Before Committing

As always, build world (or appropriate parts of it).

Check the changes with `svn diff` and `svn stat`. Make sure all the files that should have been added or deleted were in fact added or deleted.

Take a closer look at any property change (marked by a `M` in the second column of `svn stat`). Normally, no `svn:mergeinfo` properties should be anywhere except the target directory (or directories).

If something looks fishy, ask for help.

### 5.4.3.8. Committing

Make sure to commit a top level directory to have the mergeinfo included as well. Do not specify individual files on the command line. For more information about committing files in general, see the relevant section of this primer.

## 5.4.4. Vendor Imports with SVN



### Important

Please read this entire section before starting a vendor import.



### Note

Patches to vendor code fall into two categories:

- Vendor patches: these are patches that have been issued by the vendor, or that have been extracted from the vendor's version control system, which address issues which in your opinion cannot wait until the next vendor release.
- FreeBSD patches: these are patches that modify the vendor code to address FreeBSD-specific issues.

The nature of a patch dictates where it should be committed:

- Vendor patches should be committed to the vendor branch, and merged from there to head. If the patch addresses an issue in a new release that is currently being imported, it *must not* be committed along with the new release: the release must be imported and tagged first, then the patch can be applied and committed. There is no need to re-tag the vendor sources after committing the patch.
- FreeBSD patches should be committed directly to head.

### 5.4.4.1. Preparing the Tree

If importing for the first time after the switch to Subversion, flattening and cleaning up the vendor tree is necessary, as well as bootstrapping the merge history in the main tree.

#### 5.4.4.1.1. Flattening

During the conversion from CVS to Subversion, vendor branches were imported with the same layout as the main tree. This means that the pf vendor sources ended up in `vendor/pf/dist/contrib/pf`. The vendor source is best directly in `vendor/pf/dist`.

To flatten the pf tree:

```
% cd vendor/pf/dist/contrib/pf
% svn mv $(svn list) ../..
% cd ../..
% svn rm contrib
% svn propdel -R svn:mergeinfo .
% svn commit
```

The `propdel` bit is necessary because starting with 1.5, Subversion will automatically add `svn:mergeinfo` to any directory that is copied or moved. In this case, as nothing is being merged from the deleted tree, they just get in the way.

Tags may be flattened as well (3, 4, 3.5 etc.); the procedure is exactly the same, only changing `dist` to 3.5 or similar, and putting the `svn commit` off until the end of the process.

#### 5.4.4.1.2. Cleaning Up

The `dist` tree can be cleaned up as necessary. Disabling keyword expansion is recommended, as it makes no sense on unmodified vendor code and in some cases it can even be harmful. OpenSSH, for example, includes two files that originated with FreeBSD and still contain the original version tags. To do this:

```
% svn propdel svn:keywords -R .
```

```
% svn commit
```

#### 5.4.4.1.3. Bootstrapping Merge History

If importing for the first time after the switch to Subversion, bootstrap `svn:mergeinfo` on the target directory in the main tree to the revision that corresponds to the last related change to the vendor tree, prior to importing new sources:

```
% cd head/contrib/pf
% svn merge --record-only svn+ssh://svn.freebsd.org/base/ vendor/pf/
dist@180876 .
% svn commit
```

#### 5.4.4.2. Importing New Sources

With two commits—one for the import itself and one for the tag—this step can optionally be repeated for every upstream release between the last import and the current import.

##### 5.4.4.2.1. Preparing the Vendor Sources

Unlike in CVS where only the needed parts were imported into the vendor tree to avoid bloating the main tree, Subversion is able to store a full distribution in the vendor tree. So, import everything, but merge only what is required.

A `svn add` is required to add any files that were added since the last vendor import, and `svn rm` is required to remove any that were removed since. Preparing sorted lists of the contents of the vendor tree and of the sources that are about to be imported is recommended, to facilitate the process.

```
% cd vendor/pf/dist
% svn list -R | grep -v '/' | sort >../old
% cd ../pf-4.3
% find . -type f | cut -c 3- | sort >../new
```

With these two files, `comm -23 ../old ../new` will list removed files (files only in `old`), while `comm -13 ../old ../new` will list added files only in `new`.

##### 5.4.4.2.2. Importing into the Vendor Tree

Now, the sources must be copied into `dist` and the `svn add` and `svn rm` commands should be used as needed:

```
% cd vendor/pf/pf-4.3
% tar cf - . | tar xf - -C ../dist
% cd ../dist
% comm -23 ../old ../new | xargs svn rm
% comm -13 ../old ../new | xargs svn --parents add
```

If any directories were removed, they will have to be `svn rm`d manually. Nothing will break if they are not, but they will remain in the tree.

Check properties on any new files. All text files should have `svn:eol-style` set to `native`. All binary files should have `svn:mime-type` set to `application/octet-stream` unless there is a more appropriate media type. Executable files should have `svn:executable` set to `*`. No other properties should exist on any file in the tree.

Committing is now possible, however it is good practice to make sure that everything is OK by using the `svn stat` and `svn diff` commands.

#### 5.4.4.2.3. Tagging

Once committed, vendor releases should be tagged for future reference. The best and quickest way to do this is directly in the repository:

```
% svn cp svn+ssh://svn.freebsd.org/base/ vendor/pf/dist svn+ssh://svn.
freebsd.org/base/ vendor/pf/4.3
```

Once that is complete, `svn up` the working copy of `vendor/pf` to get the new tag, although this is rarely needed.

If creating the tag in the working copy of the tree, `svn:mergeinfo` results must be removed:

```
% cd vendor/pf
% svn cp dist 4.3
% svn propdel svn:mergeinfo -R 4.3
```

#### 5.4.4.3. Merging to Head

```
% cd head/contrib/pf
% svn up
% svn merge --accept=postpone svn+ssh://svn.freebsd.org/base/ vendor/
pf/dist .
```

The `--accept=postpone` tells Subversion that it should not complain because merge conflicts will be taken care of manually.



#### Tip

The `cvs2svn` changeover occurred on June 3, 2008. When performing vendor merges for packages which were already present and converted by the `cvs2svn` process, the command used to merge `/vendor/package_name /dist` to `/head/package_location` (for example, `head/contrib/sendmail`) must use `-c REV` to indicate the revision to merge from the `/vendor` tree. For example:

```
% svn checkout svn+ssh://svn.freebsd.org/base/head/
contrib/ sendmail
% cd sendmail
```



```
% svn merge -c r261190 ^/vendor/sendmail/dist .
```

^ is an alias for the repository path.



### Note

If using the Zsh shell, the ^ must be escaped with \. This means ^/head should be \^/head.

It is necessary to resolve any merge conflicts. This process is the same in SVN as in CVS.

Make sure that any files that were added or removed in the vendor tree have been properly added or removed in the main tree. To check diffs against the vendor branch:

```
% svn diff --no-diff-deleted --old=svn+ssh://svn.freebsd.org/  
base/vendor/pf/dist --new=.
```

The `--no-diff-deleted` tells Subversion not to complain about files that are in the vendor tree but not in the main tree, i.e., things that would have previously been removed before the vendor import, like for example the vendor's makefiles and configure scripts.

Using CVS, once a file was off the vendor branch, it was not able to be put back. With Subversion, there is no concept of on or off the vendor branch. If a file that previously had local modifications, to make it not show up in diffs in the vendor tree, all that has to be done is remove any left-over cruft like FreeBSD version tags, which is much easier.

If any changes are required for the world to build with the new sources, make them now, and keep testing until everything builds and runs perfectly.

## 5.4.4.4. Committing the Vendor Import

Committing is now possible! Everything must be committed in one go. If done properly, the tree will move from a consistent state with old code, to a consistent state with new code.

## 5.4.4.5. From Scratch

### 5.4.4.5.1. Importing into the Vendor Tree

This section is an example of importing and tagging byacc into head.

First, prepare the directory in vendor:

```
% svn co --depth immediates $FSVN/vendor
```

```
% cd vendor
% svn mkdir byacc
% svn mkdir byacc/dist
```

Now, import the sources into the `dist` directory. Once the files are in place, `svn add` the new ones, then `svn commit` and tag the imported version. To save time and bandwidth, direct remote committing and tagging is possible:

```
% svn cp -m "Tag byacc 20120115" $FSVN/vendor/byacc/dist $FSVN/vendor/
byacc/20120115
```

#### 5.4.4.5.2. Merging to head

Due to this being a new file, copy it for the merge:

```
% svn cp -m "Import byacc to contrib" $FSVN/vendor/byacc/dist $FSVN/
head/contrib/byacc
```

Working normally on newly imported sources is still possible.

#### 5.4.5. Reverting a Commit

Reverting a commit to a previous version is fairly easy:

```
% svn merge -r179454:179453 ROADMAP.txt
% svn commit
```

Change number syntax, with negative meaning a reverse change, can also be used:

```
% svn merge -c -179454 ROADMAP.txt
% svn commit
```

This can also be done directly in the repository:

```
% svn merge -r179454:179453 svn+ssh://svn.freebsd.org/base/ROADMAP.txt
```



#### Note

It is important to ensure that the mergeinfo is correct when reverting a file in order to permit `svn mergeinfo --eligible` to work as expected.

Reverting the deletion of a file is slightly different. Copying the version of the file that predates the deletion is required. For example, to restore a file that was deleted in revision N, restore version N-1:

```
% svn copy svn+ssh://svn.freebsd.org/base/ROADMAP.txt@179454
```

```
% svn commit
```

or, equally:

```
% svn copy svn+ssh://svn.freebsd.org/base/ROADMAP.txt@179454 svn  
+ssh://svn.freebsd.org/base
```

Do *not* simply recreate the file manually and `svn add` it—this will cause history to be lost.

### 5.4.6. Fixing Mistakes

While we can do surgery in an emergency, do not plan on having mistakes fixed behind the scenes. Plan on mistakes remaining in the logs forever. Be sure to check the output of `svn status` and `svn diff` before committing.

Mistakes will happen but, they can generally be fixed without disruption.

Take a case of adding a file in the wrong location. The right thing to do is to `svn move` the file to the correct location and commit. This causes just a couple of lines of metadata in the repository journal, and the logs are all linked up correctly.

The wrong thing to do is to delete the file and then `svn add` an independent copy in the correct location. Instead of a couple of lines of text, the repository journal grows an entire new copy of the file. This is a waste.

### 5.4.7. Setting up a svnsync Mirror

You probably do not want to do this unless there is a good reason for it. Such reasons might be to support many multiple local read-only client machines, or if your network bandwidth is limited. Starting a fresh mirror from empty would take a very long time. Expect a minimum of 10 hours for high speed connectivity. If you have international links, expect this to take 4 to 10 times longer.

A far better option is to grab a seed file. It is large (~1GB) but will consume less network traffic and take less time to fetch than a `svnsync` will. This is possible in one of the following three ways:

```
% rsync -va --partial --progress freefall:/home/peter/svnmirror-base-  
r179637.tbz2 .
```

```
% rsync -va --partial --progress rsync://repoman.freebsd.org:50873/  
svnseed/svnmirror-base-r215629.tar.xz .
```

```
% fetch ftp://ftp.freebsd.org/pub/FreeBSD/development/subversion/  
svnmirror-base-r221445.tar.xz
```

Once you have the file, extract it to somewhere like `home/svnmirror/base/`. Then, update it, so that it fetches changes since the last revision in the archive:

```
% svn sync file:///home/svnmirror/base
```

You can then set that up to run from [cron\(8\)](#), do checkouts locally, set up a svnserve server for your local machines to talk to, etc.

The seed mirror is set to fetch from `svn://svn.freebsd.org/base` . The configuration for the mirror is stored in `revprop 0` on the local mirror. To see the configuration, try:

```
% svn proplist -v --revprop -r 0 file:///home/svnmirror/base
```

Use `propset` to change things.

### 5.4.8. Committing High-ASCII Data

Files that have high-ASCII bits are considered binary files in SVN, so the pre-commit checks fail and indicate that the `mime-type` property should be set to `application/octet-stream` . However, the use of this is discouraged, so please do not set it. The best way is always avoiding high-ASCII data, so that it can be read everywhere with any text editor but if it is not avoidable, instead of changing the mime-type, set the `fbbsd:notbinary` property with `propset` :

```
% svn propset fbbsd:notbinary yes foo.data
```

### 5.4.9. Maintaining a Project Branch

A project branch is one that is synced to head (or another branch) is used to develop a project then commit it back to head. In SVN, “dolphin” branching is used for this. A “dolphin” branch is one that diverges for a while and is finally committed back to the original branch. During development code migration in one direction (from head to the branch only). No code is committed back to head until the end. Once you commit back at the end, the branch is dead (although you can have a new branch with the same name after you delete the branch if you want).

As per [http://people.freebsd.org/~peter/svn\\_notes.txt](http://people.freebsd.org/~peter/svn_notes.txt), work that is intended to be merged back into HEAD should be in `base/projects/` . If you are doing work that is beneficial to the FreeBSD community in some way but not intended to be merged directly back into HEAD then the proper location is `base/user/ your-name /` . [This page](#) contains further details.

To create a project branch:

```
% svn copy svn+ssh://svn.freebsd.org/base/head svn+ssh://svn.freebsd.org/base/projects/spif
```

To merge changes from HEAD back into the project branch:

```
% cd copy_of_spif
% svn merge svn+ssh://svn.freebsd.org/base/head
```

```
% svn commit
```

It is important to resolve any merge conflicts before committing.

## 5.5. Some Tips

In commit logs etc., “rev 179872” should be spelled “r179872” as per convention.

Speeding up svn is possible by adding the following to `~/.ssh/config` :

```
Host *
ControlPath ~/.ssh/sockets/master-%l-%r@h:%p
ControlMaster auto
ControlPersist yes
```

and then typing

```
mkdir ~/.ssh/sockets
```

Checking out a working copy with a stock Subversion client without FreeBSD-specific patches (`OPTIONS_SET=FREEBSD_TEMPLATE` ) will mean that `$FreeBSD$` tags will not be expanded. Once the correct version has been installed, trick Subversion into expanding them like so:

```
% svn propdel -R svn:keywords .
% svn revert -R .
```

This will wipe out uncommitted patches.

## 6. Setup, Conventions, and Traditions

There are a number of things to do as a new developer. The first set of steps is specific to committers only. These steps must be done by a mentor for those who are not committers.

### 6.1. For New Committers

Those who have been given commit rights to the FreeBSD repositories must follow these steps.

- Get mentor approval before committing each of these changes!
- The `.ent` and `.xml` files mentioned below exist in the FreeBSD Documentation Project SVN repository at [svn.FreeBSD.org/doc/](http://svn.FreeBSD.org/doc/) .
- New files that do not have the `FreeBSD=%H svn:keywords` property will be rejected when attempting to commit them to the repository. Be sure to read [Section 5.3.7, “Adding and Removing Files”](#) regarding adding and removing files. Verify that `~/.sub-`

version/config contains the necessary “auto-props” entries from auto-props.txt mentioned there.

- All src commits should go to FreeBSD-CURRENT first before being merged to FreeBSD-STABLE. The FreeBSD-STABLE branch must maintain ABI and API compatibility with earlier versions of that branch. Do not merge changes that break this compatibility.

#### Procedure 1. Steps for New Committers

##### 1. Add an Author Entity

doc/head/share/xml/authors.ent — Add an author entity. Later steps depend on this entity, and missing this step will cause the doc/ build to fail. This is a relatively easy task, but remains a good first test of version control skills.

##### 2. Update the List of Developers and Contributors

doc/head/en\_US.ISO8859-1/  
articles/contributors/contrib.committers.xml — Add an entry to the “Developers” section of the [Contributors List](#). Entries are sorted by last name.

doc/head/en\_US.ISO8859-1/  
articles/contributors/contrib.additional.xml — Remove the entry from the “Additional Contributors” section. Entries are sorted by last name.

##### 3. Add a News Item

doc/head/share/xml/news.xml — Add an entry. Look for the other entries that announce new committers and follow the format. Use the date from the commit bit approval email from <[core@FreeBSD.org](mailto:core@FreeBSD.org)>.

##### 4. Add a PGP Key

doc/head/share/pgpkeys/pgpkeys.ent and doc/head/share/pgpkeys/pgpkeys-developers.xml — Add your PGP or GnuPG key. Those who do not yet have a key should see [Section 2.1, “Creating a Key”](#).

Dag-Erling C. Smørgrav <[des@FreeBSD.org](mailto:des@FreeBSD.org)> has written a shell script (doc/head/share/pgpkeys/addkey.sh) to make this easier. See the [README](#) file for more information.

Use doc/head/share/pgpkeys/checkkey.sh to verify that keys meet minimal best-practices standards.

After adding and checking a key, add both updated files to source control and then commit them. Entries in this file are sorted by last name.



## Note

It is very important to have a current PGP/GnuPG key in the repository. The key may be required for positive identification of a committer. For example, the FreeBSD Administrators <[admins@FreeBSD.org](mailto:admins@FreeBSD.org)> might need it for account recovery. A complete keyring of FreeBSD.org users is available for download from <http://www.FreeBSD.org/doc/pgpkeyring.txt>.

### 5. Update Mentor and Mentee Information

`base/head/share/misc/committers- repository .dot` — Add an entry to the current committers section, where *repository* is `doc`, `ports`, or `src`, depending on the commit privileges granted.

Add an entry for each additional mentor/mentee relationship in the bottom section.

### 6. Generate a Kerberos Password

See [Section 3, “Kerberos and LDAP web Password for FreeBSD Cluster”](#) to generate or set a Kerberos for use with other FreeBSD services like the bug tracking database.

### 7. Optional: Enable Wiki Account

[FreeBSD Wiki](#) Account — A wiki account allows sharing projects and ideas. Those who do not yet have an account can contact <[clusteradm@FreeBSD.org](mailto:clusteradm@FreeBSD.org)> to obtain one.

### 8. Optional: Update Wiki Information

Wiki Information - After gaining access to the wiki, some people add entries to the [How We Got Here](#), [Irc Nicks](#), and [Dogs of FreeBSD](#) pages.

### 9. Optional: Update Ports with Personal Information

`ports/astro/xeearth/files/freebsd.committers.markers` and `src/usr.bin/calendar/calendars/calendar.freebsd` - Some people add entries for themselves to these files to show where they are located or the date of their birthday.

### 10. Optional: Prevent Duplicate Mailings

Subscribers to [svn-src-all](#), [svn-ports-all](#) or [svn-doc-all](#) might wish to unsubscribe to avoid receiving duplicate copies of commit messages and followups.

## 6.2. For Everyone

1. Introduce yourself to the other developers, otherwise no one will have any idea who you are or what you are working on. The introduction need not be a comprehensive biography, just write a paragraph or two about who you are, what you plan to be working on as a developer in FreeBSD, and who will be your mentor. Email this to the FreeBSD developers mailing list and you will be on your way!
2. Log into `hub.FreeBSD.org` and create a `/var/forward/user` (where *user* is your username) file containing the e-mail address where you want mail addressed to *yourusername*@FreeBSD.org to be forwarded. This includes all of the commit messages as well as any other mail addressed to the FreeBSD committer's mailing list and the FreeBSD developers mailing list. Really large mailboxes which have taken up permanent residence on hub often get “accidentally” truncated without warning, so forward it or read it and you will not lose it.

Due to the severe load dealing with SPAM places on the central mail servers that do the mailing list processing the front-end server does do some basic checks and will drop some messages based on these checks. At the moment proper DNS information for the connecting host is the only check in place but that may change. Some people blame these checks for bouncing valid email. If you want these checks turned off for your email you can place a file named `.spam_lover` in your home directory on `freefall.FreeBSD.org` to disable the checks for your email.



### Note

Those who are developers but not committers will not be subscribed to the committers or developers mailing lists. The subscriptions are derived from the access rights.

## 6.3. Mentors

All new developers have a mentor assigned to them for the first few months. A mentor is responsible for teaching the mentee the rules and conventions of the project and guiding their first steps in the developer community. The mentor is also personally responsible for the mentee's actions during this initial period.

For committers: do not commit anything without first getting mentor approval. Document that approval with an `Approved by:` line in the commit message.

When the mentor decides that a mentee has learned the ropes and is ready to commit on their own, the mentor announces it with a commit to `mentors`.



## 7. Commit Log Messages

This section contains some suggestions and traditions for how commit logs are formatted.

As well as including an informative message with each commit you may need to include some additional information.

This information consists of one or more lines containing the key word or phrase, a colon, tabs for formatting, and then the additional information.

The key words or phrases are:

PR:	The problem report (if any) which is affected (typically, by being closed) by this commit. Only include one PR per line as the automated scripts which parse this line cannot understand more than one.
Submitted by:	<p>The name and e-mail address of the person that submitted the fix; for developers, just the username on the FreeBSD cluster.</p> <p>If the submitter is the maintainer of the port to which you are committing, include "(maintainer)" after the email address.</p> <p>Avoid obfuscating the email address of the submitter as this adds additional work when searching logs.</p>
Reviewed by:	The name and e-mail address of the person or people that reviewed the change; for developers, just the username on the FreeBSD cluster. If a patch was submitted to a mailing list for review, and the review was favorable, then just include the list name.
Approved by:	The name and e-mail address of the person or people that approved the change; for developers, just the username on the FreeBSD cluster. It is customary to get prior approval for a commit if it is to an area of the tree to which you do not usually commit. In addition, during the run up to a new release all commits <i>must</i> be approved by the release engineering team.

	<p>If these are your first commits then you should have passed them past your mentor first, and you should list your mentor, as in <code>``username-of-mentor (mentor)``</code>.</p> <p>If a team approved these commits then include the team name followed by the username of the approver in parentheses. For example: <code>``re@ (username)``</code></p>
Obtained from:	The name of the project (if any) from which the code was obtained. Do not use this line for the name of an individual person.
MFC after:	If you wish to receive an e-mail reminder to MFC at a later date, specify the number of days, weeks, or months after which an MFC is planned.
Relnotes:	If the change is a candidate for inclusion in the release notes for the next release from the branch, set to yes.
Security:	If the change is related to a security vulnerability or security exposure, include one or more references or a description of the issue. If possible, include a VuXML URL or a CVE ID.

### Example 1. Commit Log for a Commit Based on a PR

You want to commit a change based on a PR submitted by John Smith containing a patch. The end of the commit message should look something like this.

...

```
PR: 12345
Submitted by: John Smith <John.Smith@example.com>
```

### Example 2. Commit Log for a Commit Needing Review

You want to change the virtual memory system. You have posted patches to the appropriate mailing list (in this case, `freebsd-arch`) and the changes have been approved.

...

Reviewed by:        `-arch`

### Example 3. Commit Log for a Commit Needing Approval

You want to commit a port. You have collaborated with the listed MAINTAINER, who has told you to go ahead and commit.

...

Approved by:        `abc (maintainer)`

Where `abc` is the account name of the person who approved.

### Example 4. Commit Log for a Commit Bringing in Code from OpenBSD

You want to commit some code based on work done in the OpenBSD project.

...

Obtained from:        `OpenBSD`

### Example 5. Commit Log for a Change to FreeBSD-CURRENT with a Planned Commit to FreeBSD-STABLE to Follow at a Later Date.

You want to commit some code which will be merged from FreeBSD-CURRENT into the FreeBSD-STABLE branch after two weeks.

```
...
```

```
MFC after:      2 weeks
```

Where 2 is the number of days, weeks, or months after which an MFC is planned. The *weeks* option may be *day, days, week, weeks, month, months*.

In many cases you may need to combine some of these.

Consider the situation where a user has submitted a PR containing code from the NetBSD project. You are looking at the PR, but it is not an area of the tree you normally work in, so you have decided to get the change reviewed by the arch mailing list. Since the change is complex, you opt to MFC after one month to allow adequate testing.

The extra information to include in the commit would look something like

### Example 6. Example Combined Commit Log

```
PR:                54321
Submitted by:      John Smith <John.Smith@example.com>
Reviewed by:       -arch
Obtained from:     NetBSD
MFC after:         1 month
Relnotes:          yes
```

## 8. Preferred License for New Files

Currently the FreeBSD Project suggests and uses the following text as the preferred license scheme:

```
/*-
```

```
* Copyright (c) [year] [your name]
* All rights reserved.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
* 1. Redistributions of source code must retain the above copyright
*    notice, this list of conditions and the following disclaimer.
* 2. Redistributions in binary form must reproduce the above copyright
*    notice, this list of conditions and the following disclaimer in the
*    documentation and/or other materials provided with the distribution.
*
* THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS ``AS IS'' AND
* ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
* PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE
* LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
* CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
* SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS
* INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN
* CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING
* IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE
* POSSIBILITY OF SUCH DAMAGE.
*
* [id for your version control system, if any]
*/
```

The FreeBSD project strongly discourages the so-called "advertising clause" in new code. Due to the large number of contributors to the FreeBSD project, complying with this clause for many commercial vendors has become difficult. If you have code in the tree with the advertising clause, please consider removing it. In fact, please consider using the above license for your code.

The FreeBSD project discourages completely new licenses and variations on the standard licenses. New licenses require the approval of the Core Team <[core@FreeBSD.org](mailto:core@FreeBSD.org)> to reside in the main repository. The more different licenses that are used in the tree, the more problems that this causes to those wishing to utilize this code, typically from unintended consequences from a poorly worded license.

Project policy dictates that code under some non-BSD licenses must be placed only in specific sections of the repository, and in some cases, compilation must be conditional or even disabled by default. For example, the GENERIC kernel must be compiled under only licenses identical to or substantially similar to the BSD license. GPL, APSL, CDDL, etc, licensed software must not be compiled into GENERIC.

Developers are reminded that in open source, getting "open" right is just as important as getting "source" right, as improper handling of intellectual property has serious consequences. Any questions or concerns should immediately be brought to the attention of the core team.

## 9. Keeping Track of Licenses Granted to the FreeBSD Project

Various software or data exist in the repositories where the FreeBSD project has been granted a special licence to be able to use them. A case in point are the Terminus fonts for use with [vt\(4\)](#). Here the author Dimitar Zhekov has allowed us to use the "Terminus BSD Console" font under a 2-clause BSD license rather than the regular Open Font License he normally uses.

It is clearly sensible to keep a record of any such license grants. To that end, the Core Team <[core@FreeBSD.org](mailto:core@FreeBSD.org)> has decided to keep an archive of them. Whenever the FreeBSD project is granted a special license we require the Core Team <[core@FreeBSD.org](mailto:core@FreeBSD.org)> to be notified. Any developers involved in arranging such a license grant, please send details to the Core Team <[core@FreeBSD.org](mailto:core@FreeBSD.org)> including:

- Contact details for people or organizations granting the special license.
- What files, directories etc. in the repositories are covered by the license grant including the revision numbers where any specially licensed material was committed.
- The date the license comes into effect from. Unless otherwise agreed, this will be the date the license was issued by the authors of the software in question.
- The license text.
- A note of any restrictions, limitations or exceptions that apply specifically to FreeBSD's usage of the licensed material.
- Any other relevant information.

Once the Core Team <[core@FreeBSD.org](mailto:core@FreeBSD.org)> is satisfied that all the necessary details have been gathered and are correct, the secretary will send a PGP-signed acknowledgement of receipt including the license details. This receipt will be persistently archived and serve as our permanent record of the license grant.

The license archive should contain only details of license grants; this is not the place for any discussions around licensing or other subjects. Access to data within the license archive will be available on request to the Core Team <[core@FreeBSD.org](mailto:core@FreeBSD.org)>.

## 10. Developer Relations

If you are working directly on your own code or on code which is already well established as your responsibility, then there is probably little need to check with other committers before jumping in with a commit. If you see a bug in an area of the system which is clearly orphaned (and there are a few such areas, to our shame), the same applies. If, however, you are about to modify something which is clearly being actively maintained by someone else (and it is only by watching the *repository*-committers mailing list that you can really get a feel for just what is and is not) then consider sending the change to them instead, just as you would have before becoming a committer. For ports, you should contact the listed MAINTAINER in the Makefile. For other parts of the repository, if you are unsure who the active maintainer might be, it may help to scan the revision history to see who has committed changes in the past. An example script that lists each person who has committed to a given file along with the number of commits each person has made can be found at on freefall at ~eadler/bin/whodid . If your queries go unanswered or the committer otherwise indicates a lack of interest in the area affected, go ahead and commit it.



### Note

Avoid sending private emails to maintainers. Other people might be interested in the conversation, not just the final output.

If you are unsure about a commit for any reason at all, have it reviewed by -hackers before committing. Better to have it flamed then and there rather than when it is part of the repository. If you do happen to commit something which results in controversy erupting, you may also wish to consider backing the change out again until the matter is settled. Remember – with a version control system we can always change it back.

Do not impugn the intentions of someone you disagree with. If they see a different solution to a problem than you, or even a different problem, it is not because they are stupid, because they have questionable parentage, or because they are trying to destroy your hard work, personal image, or FreeBSD, but simply because they have a different outlook on the world. Different is good.

Disagree honestly. Argue your position from its merits, be honest about any shortcomings it may have, and be open to seeing their solution, or even their vision of the problem, with an open mind.

Accept correction. We are all fallible. When you have made a mistake, apologize and get on with life. Do not beat up yourself, and certainly do not beat up others for your mistake. Do not waste time on embarrassment or recrimination, just fix the problem and move on.

Ask for help. Seek out (and give) peer reviews. One of the ways open source software is supposed to excel is in the number of eyeballs applied to it; this does not apply if nobody will review code.

## 11. If in Doubt...

When you are not sure about something, whether it be a technical issue or a project convention be sure to ask. If you stay silent you will never make progress.

If it relates to a technical issue ask on the public mailing lists. Avoid the temptation to email the individual person that knows the answer. This way everyone will be able to learn from the question and the answer.

For project specific or administrative questions you should ask, in order:

- Your mentor or former mentor.
- An experienced committer on IRC, email, etc.
- Any team with a "hat", as they should give you a definitive answer.
- If still not sure, ask on FreeBSD developers mailing list.

Once your question is answered, if no one pointed you to documentation that spelled out the answer to your question, document it, as others will have the same question.

## 12. Bugzilla

The FreeBSD Project utilizes Bugzilla for tracking bugs and change requests. Be sure that if you commit a fix or suggestion found in the PR database to close it. It is also considered nice if you take time to close any PRs associated with your commits, if appropriate.

Committers with non-FreeBSD.org Bugzilla accounts can have the old account merged with the FreeBSD.org account by entering a new bug. Choose Supporting Services as the Product, and Bug Tracker as the Component.

You can find out more about Bugzilla at:



- [FreeBSD Problem Report Handling Guidelines](#)
- <http://www.FreeBSD.org/support.html>

## 13. Who's Who

Besides the repository meisters, there are other FreeBSD project members and teams whom you will probably get to know in your role as a committer. Briefly, and by no means all-inclusively, these are:

Documentation Engineering Team <[doceng@FreeBSD.org](mailto:doceng@FreeBSD.org)>

doceng is the group responsible for the documentation build infrastructure, approving new documentation committers, and ensuring that the FreeBSD website and documentation on the FTP site is up to date with respect to the CVS tree. It is not a conflict resolution body. The vast majority of documentation related discussion takes place on the [FreeBSD documentation project mailing list](#). More details regarding the doceng team can be found in its [charter](#). Committers interested in contributing to the documentation should familiarize themselves with the [Documentation Project Primer](#).

Bruce Evans <[bde@FreeBSD.org](mailto:bde@FreeBSD.org)>

Bruce is the Style Police-Meister. When you do a commit that could have been done better, Bruce will be there to tell you. Be thankful that someone is. Bruce is also very knowledgeable on the various standards applicable to FreeBSD.

Glen Barber <[gjb@FreeBSD.org](mailto:gjb@FreeBSD.org)>, Konstantin Belousov <[kib@FreeBSD.org](mailto:kib@FreeBSD.org)>, Joel Dahl <[joel@FreeBSD.org](mailto:joel@FreeBSD.org)>, Marc Fonvieille <[blackend@FreeBSD.org](mailto:blackend@FreeBSD.org)>, Steven Kreuzer <[skreuzer@FreeBSD.org](mailto:skreuzer@FreeBSD.org)>, Xin Li <[delphij@FreeBSD.org](mailto:delphij@FreeBSD.org)>, Josh Paetzel <[jpaetzel@FreeBSD.org](mailto:jpaetzel@FreeBSD.org)>, Craig Rodrigues <[rodrigc@FreeBSD.org](mailto:rodrigc@FreeBSD.org)>, Hiroki Sato <[hros@FreeBSD.org](mailto:hros@FreeBSD.org)>, Gleb Smirnov <[glebius@FreeBSD.org](mailto:glebius@FreeBSD.org)>, Ken Smith <[kensmith@FreeBSD.org](mailto:kensmith@FreeBSD.org)>, Marius Strobl <[marius@FreeBSD.org](mailto:marius@FreeBSD.org)>

These are the members of the Release Engineering Team <[re@FreeBSD.org](mailto:re@FreeBSD.org)>. This team is responsible for setting release deadlines and controlling the release process. During code freezes, the release engineers have final authority on all changes to the system for whichever branch is pending release status. If there is something you want merged from FreeBSD-CURRENT to FreeBSD-STABLE (whatever values those may have at any given time), these are the people to talk to about it.

Hiroki is also the keeper of the release documentation (src/release/doc/\*). If you commit a change that you think is worthy of mention in the release notes, please make sure he knows about it. Better still, send him a patch with your suggested commentary.

Dag-Erling C. Smørgrav <[des@FreeBSD.org](mailto:des@FreeBSD.org)>

Dag-Erling is the [FreeBSD Security Officer](#) and oversees the Security Officer Team <[security-officer@FreeBSD.org](mailto:security-officer@FreeBSD.org)>.

Garrett Wollman <[wollman@FreeBSD.org](mailto:wollman@FreeBSD.org)>

If you need advice on obscure network internals or are not sure of some potential change to the networking subsystem you have in mind, Garrett is someone to talk to. Garrett is also very knowledgeable on the various standards applicable to FreeBSD.

FreeBSD committer's mailing list

[svn-src-all](#), [svn-ports-all](#) and [svn-doc-all](#) are the mailing lists that the version control system uses to send commit messages to. You should *never* send email directly to these lists. You should only send replies to this list when they are short and are directly related to a commit.

FreeBSD developers mailing list

All committers are subscribed to -developers. This list was created to be a forum for the committers “community” issues. Examples are Core voting, announcements, etc.

The FreeBSD developers mailing list is for the exclusive use of FreeBSD committers. In order to develop FreeBSD, committers must have the ability to openly discuss matters that will be resolved before they are publicly announced. Frank discussions of work in progress are not suitable for open publication and may harm FreeBSD.

All FreeBSD committers are expected not to not publish or forward messages from the FreeBSD developers mailing list outside the list membership without permission of all of the authors. Violators will be removed from the FreeBSD developers mailing list, resulting in a suspension of commit privileges. Repeated or flagrant violations may result in permanent revocation of commit privileges.

This list is *not* intended as a place for code reviews or for any technical discussion. In fact using it as such hurts the FreeBSD Project as it gives a sense of a closed list where general decisions affecting all of the FreeBSD using community are made without being “open”. Last, but not least *never, never ever, email the FreeBSD developers mailing list and CC:/BCC: another FreeBSD list*. Never, ever email another FreeBSD email list and CC:/BCC: the FreeBSD developers mailing list. Doing so can greatly diminish the benefits of this list.

## 14. SSH Quick-Start Guide

1. If you do not wish to type your password in every time you use [ssh\(1\)](#), and you use RSA or DSA keys to authenticate, [ssh-agent\(1\)](#) is there for your convenience. If you want to use [ssh-agent\(1\)](#), make sure that you run it before running other applications. X users, for example, usually do this from their `.xsession` or `.xinitrc`. See [ssh-agent\(1\)](#) for details.

2. Generate a key pair using [ssh-keygen\(1\)](#). The key pair will wind up in your `$HOME/.ssh/` directory.
3. Send your public key (`$HOME/.ssh/id_dsa.pub` or `$HOME/.ssh/id_rsa.pub`) to the person setting you up as a committer so it can be put into *yourlogin* in `/etc/ssh-keys/` on *freefall*.

Now you should be able to use [ssh-add\(1\)](#) for authentication once per session. This will prompt you for your private key's pass phrase, and then store it in your authentication agent ([ssh-agent\(1\)](#)). If you no longer wish to have your key stored in the agent, issuing `ssh-add -d` will remove it.

Test by doing something such as `ssh freefall.FreeBSD.org ls /usr`.

For more information, see [security/openssh](#), [ssh\(1\)](#), [ssh-add\(1\)](#), [ssh-agent\(1\)](#), [ssh-keygen\(1\)](#), and [scp\(1\)](#).

For information on adding, changing, or removing [ssh\(1\)](#) keys, see this [article](#).

## 15. Coverity® Availability for FreeBSD Committers

All FreeBSD developers can obtain access to Coverity analysis results of all FreeBSD Project software. All who are interested in obtaining access to the analysis results of the automated Coverity runs, can sign up at [Coverity Scan](#).

The FreeBSD wiki includes a mini-guide for developers who are interested in working with the Coverity® analysis reports: <http://wiki.freebsd.org/CoverityPrevent>. Please note that this mini-guide is only readable by FreeBSD developers, so if you cannot access this page, you will have to ask someone to add you to the appropriate Wiki access list.

Finally, all FreeBSD developers who are going to use Coverity® are always encouraged to ask for more details and usage information, by posting any questions to the mailing list of the FreeBSD developers.

## 16. The FreeBSD Committers' Big List of Rules

1. Respect other committers.
2. Respect other contributors.
3. Discuss any significant change *before* committing.
4. Respect existing maintainers (if listed in the MAINTAINER field in `Makefile` or in `MAINTAINER` in the top-level directory).

5. Any disputed change must be backed out pending resolution of the dispute if requested by a maintainer. Security related changes may override a maintainer's wishes at the Security Officer's discretion.
6. Changes go to FreeBSD-CURRENT before FreeBSD-STABLE unless specifically permitted by the release engineer or unless they are not applicable to FreeBSD-CURRENT. Any non-trivial or non-urgent change which is applicable should also be allowed to sit in FreeBSD-CURRENT for at least 3 days before merging so that it can be given sufficient testing. The release engineer has the same authority over the FreeBSD-STABLE branch as outlined for the maintainer in rule #5.
7. Do not fight in public with other committers; it looks bad. If you must “strongly disagree” about something, do so only in private.
8. Respect all code freezes and read the committers and developers mailing lists in a timely manner so you know when a code freeze is in effect.
9. When in doubt on any procedure, ask first!
10. Test your changes before committing them.
11. Do not commit to anything under the `src/contrib`, `src/crypto`, or `src/sys/contrib` trees without *explicit* approval from the respective maintainer(s).

As noted, breaking some of these rules can be grounds for suspension or, upon repeated offense, permanent removal of commit privileges. Individual members of core have the power to temporarily suspend commit privileges until core as a whole has the chance to review the issue. In case of an “emergency” (a committer doing damage to the repository), a temporary suspension may also be done by the repository meisters. Only a 2/3 majority of core has the authority to suspend commit privileges for longer than a week or to remove them permanently. This rule does not exist to set core up as a bunch of cruel dictators who can dispose of committers as casually as empty soda cans, but to give the project a kind of safety fuse. If someone is out of control, it is important to be able to deal with this immediately rather than be paralyzed by debate. In all cases, a committer whose privileges are suspended or revoked is entitled to a “hearing” by core, the total duration of the suspension being determined at that time. A committer whose privileges are suspended may also request a review of the decision after 30 days and every 30 days thereafter (unless the total suspension period is less than 30 days). A committer whose privileges have been revoked entirely may request a review after a period of 6 months has elapsed. This review policy is *strictly informal* and, in all cases, core reserves the right to either act on or disregard requests for review if they feel their original decision to be the right one.

In all other aspects of project operation, core is a subset of committers and is bound by the *same rules*. Just because someone is in core this does not mean that they have special dispensation to step outside any of the lines painted here; core's “special powers” only

kick in when it acts as a group, not on an individual basis. As individuals, the core team members are all committers first and core second.

## 16.1. Details

### 1. Respect other committers.

This means that you need to treat other committers as the peer-group developers that they are. Despite our occasional attempts to prove the contrary, one does not get to be a committer by being stupid and nothing rankles more than being treated that way by one of your peers. Whether we always feel respect for one another or not (and everyone has off days), we still have to *treat* other committers with respect at all times, on public forums and in private email.

Being able to work together long term is this project's greatest asset, one far more important than any set of changes to the code, and turning arguments about code into issues that affect our long-term ability to work harmoniously together is just not worth the trade-off by any conceivable stretch of the imagination.

To comply with this rule, do not send email when you are angry or otherwise behave in a manner which is likely to strike others as needlessly confrontational. First calm down, then think about how to communicate in the most effective fashion for convincing the other person(s) that your side of the argument is correct, do not just blow off some steam so you can feel better in the short term at the cost of a long-term flame war. Not only is this very bad “energy economics”, but repeated displays of public aggression which impair our ability to work well together will be dealt with severely by the project leadership and may result in suspension or termination of your commit privileges. The project leadership will take into account both public and private communications brought before it. It will not seek the disclosure of private communications, but it will take it into account if it is volunteered by the committers involved in the complaint.

All of this is never an option which the project's leadership enjoys in the slightest, but unity comes first. No amount of code or good advice is worth trading that away.

### 2. Respect other contributors.

You were not always a committer. At one time you were a contributor. Remember that at all times. Remember what it was like trying to get help and attention. Do not forget that your work as a contributor was very important to you. Remember what it was like. Do not discourage, belittle, or demean contributors. Treat them with respect. They are our committers in waiting. They are every bit as important to the project as committers. Their contributions are as valid and as important as your own. After all, you made many contributions before you became a committer. Always remember that.

Consider the points raised under [1](#) and apply them also to contributors.

---

3. Discuss any significant change *before* committing.

The repository is not where changes should be initially submitted for correctness or argued over, that should happen first in the mailing lists and the commit should only happen once something resembling consensus has been reached. This does not mean that you have to ask permission before correcting every obvious syntax error or manual page misspelling, simply that you should try to develop a feel for when a proposed change is not quite such a no-brainer and requires some feedback first. People really do not mind sweeping changes if the result is something clearly better than what they had before, they just do not like being *surprized* by those changes. The very best way of making sure that you are on the right track is to have your code reviewed by one or more other committers.

When in doubt, ask for review!

4. Respect existing maintainers if listed.

Many parts of FreeBSD are not “owned” in the sense that any specific individual will jump up and yell if you commit a change to “their” area, but it still pays to check first. One convention we use is to put a maintainer line in the Makefile for any package or subtree which is being actively maintained by one or more people; see [http://www.FreeBSD.org/doc/en\\_US.ISO8859-1/books/developers-hand-book/policies.html](http://www.FreeBSD.org/doc/en_US.ISO8859-1/books/developers-hand-book/policies.html) for documentation on this. Where sections of code have several maintainers, commits to affected areas by one maintainer need to be reviewed by at least one other maintainer. In cases where the “maintainer-ship” of something is not clear, you can also look at the repository logs for the file(s) in question and see if someone has been working recently or predominantly in that area.

Other areas of FreeBSD fall under the control of someone who manages an overall category of FreeBSD evolution, such as internationalization or networking. See <http://www.FreeBSD.org/administration.html> for more information on this.

5. Any disputed change must be backed out pending resolution of the dispute if requested by a maintainer. Security related changes may override a maintainer's wishes at the Security Officer's discretion.

This may be hard to swallow in times of conflict (when each side is convinced that they are in the right, of course) but a version control system makes it unnecessary to have an ongoing dispute raging when it is far easier to simply reverse the disputed change, get everyone calmed down again and then try to figure out what is the best way to proceed. If the change turns out to be the best thing after all, it can be easily brought back. If it turns out not to be, then the users did not have to live with the bogus change in the tree while everyone was busily debating its merits. People very rarely call for back-outs in the repository since discussion generally exposes bad or controversial changes before the commit even happens, but on such rare occasions

the back-out should be done without argument so that we can get immediately on to the topic of figuring out whether it was bogus or not.

6. Changes go to FreeBSD-CURRENT before FreeBSD-STABLE unless specifically permitted by the release engineer or unless they are not applicable to FreeBSD-CURRENT. Any non-trivial or non-urgent change which is applicable should also be allowed to sit in FreeBSD-CURRENT for at least 3 days before merging so that it can be given sufficient testing. The release engineer has the same authority over the FreeBSD-STABLE branch as outlined in rule #5.

This is another “do not argue about it” issue since it is the release engineer who is ultimately responsible (and gets beaten up) if a change turns out to be bad. Please respect this and give the release engineer your full cooperation when it comes to the FreeBSD-STABLE branch. The management of FreeBSD-STABLE may frequently seem to be overly conservative to the casual observer, but also bear in mind the fact that conservatism is supposed to be the hallmark of FreeBSD-STABLE and different rules apply there than in FreeBSD-CURRENT. There is also really no point in having FreeBSD-CURRENT be a testing ground if changes are merged over to FreeBSD-STABLE immediately. Changes need a chance to be tested by the FreeBSD-CURRENT developers, so allow some time to elapse before merging unless the FreeBSD-STABLE fix is critical, time sensitive or so obvious as to make further testing unnecessary (spelling fixes to manual pages, obvious bug/typo fixes, etc.) In other words, apply common sense.

Changes to the security branches (for example, releng/9.3) must be approved by a member of the Security Officer Team <[security-officer@FreeBSD.org](mailto:security-officer@FreeBSD.org)>, or in some cases, by a member of the Release Engineering Team <[re@FreeBSD.org](mailto:re@FreeBSD.org)>.

7. Do not fight in public with other committers; it looks bad. If you must “strongly disagree” about something, do so only in private.

This project has a public image to uphold and that image is very important to all of us, especially if we are to continue to attract new members. There will be occasions when, despite everyone's very best attempts at self-control, tempers are lost and angry words are exchanged. The best thing that can be done in such cases is to minimize the effects of this until everyone has cooled back down. That means that you should not air your angry words in public and you should not forward private correspondence to public mailing lists or aliases. What people say one-to-one is often much less sugar-coated than what they would say in public, and such communications therefore have no place there - they only serve to inflame an already bad situation. If the person sending you a flame-o-gram at least had the grace to send it privately, then have the grace to keep it private yourself. If you feel you are being unfairly treated by another developer, and it is causing you anguish, bring the matter up with core rather than taking it public. Core will do its best to play peace makers and get things back to sanity. In cases where the dispute involves a change to the codebase and the participants do not appear to be reaching an amicable agreement, core may appoint a mutually-agreeable third party

to resolve the dispute. All parties involved must then agree to be bound by the decision reached by this third party.

8. Respect all code freezes and read the committers and developers mailing list on a timely basis so you know when a code freeze is in effect.

Committing unapproved changes during a code freeze is a really big mistake and committers are expected to keep up-to-date on what is going on before jumping in after a long absence and committing 10 megabytes worth of accumulated stuff. People who abuse this on a regular basis will have their commit privileges suspended until they get back from the FreeBSD Happy Reeducation Camp we run in Greenland.

9. When in doubt on any procedure, ask first!

Many mistakes are made because someone is in a hurry and just assumes they know the right way of doing something. If you have not done it before, chances are good that you do not actually know the way we do things and really need to ask first or you are going to completely embarrass yourself in public. There is no shame in asking “how in the heck do I do this?” We already know you are an intelligent person; otherwise, you would not be a committer.

10. Test your changes before committing them.

This may sound obvious, but if it really were so obvious then we probably would not see so many cases of people clearly not doing this. If your changes are to the kernel, make sure you can still compile both GENERIC and LINT. If your changes are anywhere else, make sure you can still make world. If your changes are to a branch, make sure your testing occurs with a machine which is running that code. If you have a change which also may break another architecture, be sure and test on all supported architectures. Please refer to the [FreeBSD Internal Page](#) for a list of available resources. As other architectures are added to the FreeBSD supported platforms list, the appropriate shared testing resources will be made available.

11. Do not commit to anything under the `src/contrib`, `src/crypto`, and `src/sys/contrib` trees without *explicit* approval from the respective maintainer(s).

The trees mentioned above are for contributed software usually imported onto a vendor branch. Committing something there, even if it does not take the file off the vendor branch, may cause unnecessary headaches for those responsible for maintaining that particular piece of software. Thus, unless you have *explicit* approval from the maintainer (or you are the maintainer), do *not* commit there!

Please note that this does not mean you should not try to improve the software in question; you are still more than welcome to do so. Ideally, you should submit your patches to the vendor. If your changes are FreeBSD-specific, talk to the maintainer; they may be willing to apply them locally. But whatever you do, do *not* commit there by yourself!



Contact the Core Team <[core@FreeBSD.org](mailto:core@FreeBSD.org)> if you wish to take up maintainership of an unmaintained part of the tree.

## 16.2. Policy on Multiple Architectures

FreeBSD has added several new architecture ports during recent release cycles and is truly no longer an i386™ centric operating system. In an effort to make it easier to keep FreeBSD portable across the platforms we support, core has developed the following mandate:

Our 32-bit reference platform is i386, and our 64-bit reference platform is sparc64. Major design work (including major API and ABI changes) must prove itself on at least one 32-bit and at least one 64-bit platform, preferably the primary reference platforms, before it may be committed to the source tree.

The i386 and sparc64 platforms were chosen due to being more readily available to developers and as representatives of more diverse processor and system designs - big versus little endian, register file versus register stack, different DMA and cache implementations, hardware page tables versus software TLB management etc.

The ia64 platform has many of the same complications that sparc64 has, but is still limited in availability to developers.

We will continue to re-evaluate this policy as cost and availability of the 64-bit platforms change.

Developers should also be aware of our Tier Policy for the long term support of hardware architectures. The rules here are intended to provide guidance during the development process, and are distinct from the requirements for features and architectures listed in that section. The Tier rules for feature support on architectures at release-time are more strict than the rules for changes during the development process.

## 16.3. Other Suggestions

When committing documentation changes, use a spell checker before committing. For all XML docs, you should also verify that your formatting directives are correct by running `make lint`.

For all on-line manual pages, run `manck` (from ports) over the manual page to verify all of the cross references and file references are correct and that the man page has all of the appropriate `MLINKs` installed.

Do not mix style fixes with new functionality. A style fix is any change which does not modify the functionality of the code. Mixing the changes obfuscates the functionality change when asking for differences between revisions, which can hide any new bugs. Do

not include whitespace changes with content changes in commits to `doc/` or `www/`. The extra clutter in the diffs makes the translators' job much more difficult. Instead, make any style or whitespace changes in separate commits that are clearly labeled as such in the commit message.

## 16.4. Deprecating Features

When it is necessary to remove functionality from software in the base system the following guidelines should be followed whenever possible:

1. Mention is made in the manual page and possibly the release notes that the option, utility, or interface is deprecated. Use of the deprecated feature generates a warning.
2. The option, utility, or interface is preserved until the next major (point zero) release.
3. The option, utility, or interface is removed and no longer documented. It is now obsolete. It is also generally a good idea to note its removal in the release notes.

## 17. Support for Multiple Architectures

FreeBSD is a highly portable operating system intended to function on many different types of hardware architectures. Maintaining clean separation of Machine Dependent (MD) and Machine Independent (MI) code, as well as minimizing MD code, is an important part of our strategy to remain agile with regards to current hardware trends. Each new hardware architecture supported by FreeBSD adds substantially to the cost of code maintenance, toolchain support, and release engineering. It also dramatically increases the cost of effective testing of kernel changes. As such, there is strong motivation to differentiate between classes of support for various architectures while remaining strong in a few key architectures that are seen as the FreeBSD “target audience”.

### 17.1. Statement of General Intent

The FreeBSD Project targets "production quality commercial off-the-shelf (COTS) workstation, server, and high-end embedded systems". By retaining a focus on a narrow set of architectures of interest in these environments, the FreeBSD Project is able to maintain high levels of quality, stability, and performance, as well as minimize the load on various support teams on the project, such as the ports team, documentation team, security officer, and release engineering teams. Diversity in hardware support broadens the options for FreeBSD consumers by offering new features and usage opportunities (such as support for 64-bit CPUs, use in embedded environments, etc.), but these benefits must always be carefully considered in terms of the real-world maintenance cost associated with additional platform support.

The FreeBSD Project differentiates platform targets into four tiers. Each tier includes a specification of the requirements for an architecture to be in that tier, as well as specifying

the obligations of developers with regards to the platform. In addition, a policy is defined regarding the circumstances required to change the tier of an architecture.

## **17.2. Tier 1: Fully Supported Architectures**

Tier 1 platforms are fully supported by the security officer, release engineering, and toolchain maintenance staff. New features added to the operating system must be fully functional across all Tier 1 architectures for every release (features which are inherently architecture-specific, such as support for hardware device drivers, may be exempt from this requirement). In general, all Tier 1 platforms must have build and Tinderbox support either in the FreeBSD.org cluster, or be easily available for all developers. Embedded platforms may substitute an emulator available in the FreeBSD cluster for actual hardware.

Tier 1 architectures are expected to be Production Quality with respects to all aspects of the FreeBSD operating system, including installation and development environments.

Tier 1 architectures are expected to be completely integrated into the source tree and have all features necessary to produce an entire system relevant for that target architecture. Tier 1 architectures generally have at least 6 active developers.

Tier 1 architectures are expected to be fully supported by the ports system. All the ports should build on a Tier 1 platform, or have the appropriate filters to prevent the inappropriate ones from building there. The packaging system must support all Tier 1 architectures. To ensure an architecture's Tier 1 status, proponents of that architecture must show that all relevant packages can be built on that platform.

Tier 1 embedded architectures must be able to cross-build packages on at least one other Tier 1 architecture. The packages must be the most relevant for the platform, but may be a non-empty subset of those that build natively.

Tier 1 architectures must be fully documented. All basic operations need to be covered by the handbook or other documents. All relevant integration documentation must also be integrated into the tree, or readily available.

Current Tier 1 platforms are i386 and amd64.

## **17.3. Tier 2: Developmental Architectures**

Tier 2 platforms are not supported by the security officer and release engineering teams. Platform maintainers are responsible for toolchain support in the tree. The toolchain maintainer is expected to work with the platform maintainers to refine these changes. Major new toolchain components are allowed to break support for Tier 2 architectures if the FreeBSD-local changes have not been incorporated upstream. The toolchain maintainers are expected to provide prompt review of any proposed changes and cannot block, through their inaction, changes going into the tree. New features added to FreeBSD should be feasible to implement on these platforms, but an implementation is not required before

the feature may be added to the FreeBSD source tree. New features that may be difficult to implement on Tier 2 architectures should provide a means of disabling them on those architectures. The implementation of a Tier 2 architecture may be committed to the main FreeBSD tree as long as it does not interfere with production work on Tier 1 platforms, or substantially with other Tier 2 platforms. Before a Tier 2 platform can be added to the FreeBSD base source tree, the platform must be able to boot multi-user on actual hardware. Generally, there must be at least three active developers working on the platform.

Tier 2 architectures are usually systems targeted at Tier 1 support, but that are still under development. Architectures reaching end of life may also be moved from Tier 1 status to Tier 2 status as the availability of resources to continue to maintain the system in a Production Quality state diminishes. Well supported niche architectures may also be Tier 2.

Tier 2 architectures may have some support for them integrated into the ports infrastructure. They may have cross compilation support added, at the discretion of portmgr. Some ports must be built natively into packages if the package system supports that architecture. If not integrated into the base system, some external patches for the architecture for ports must be available.

Tier 2 architectures can be integrated into the FreeBSD handbook. The basics for how to get a system running must be documented, although not necessarily for every single board or system a Tier 2 architecture supports. The supported hardware list must exist and should be no more than a couple of months old. It should be integrated into the FreeBSD documentation.

Current Tier 2 platforms are arm, ia64, pc98, powerpc, and sparc64.

## **17.4. Tier 3: Experimental Architectures**

Tier 3 platforms are not supported by the security officer and release engineering teams. At the discretion of the toolchain maintainer, they may be supported in the toolchain. Tier 3 platforms are architectures in the early stages of development, for non-mainstream hardware platforms, or which are considered legacy systems unlikely to see broad future use. New Tier 3 systems will not be committed to the base source tree. Support for Tier 3 systems may be worked on in the FreeBSD Perforce Repository, providing source control and easier change integration from the main FreeBSD tree. Platforms that transition to Tier 3 status may be removed from the tree if they are no longer actively supported by the FreeBSD developer community at the discretion of the release engineer.

Tier 3 platforms may have ports support, either integrated or external, but do not require it.

Tier 3 platforms must have the basics documented for how to build a kernel and how to boot it on at least one target hardware or emulation environment. This documentation need not be integrated into the FreeBSD tree.

Current Tier 3 platforms are mips and S/390®.

## 17.5. Tier 4: Unsupported Architectures

Tier 4 systems are not supported in any form by the project.

All systems not otherwise classified into a support tier are Tier 4 systems.

## 17.6. Policy on Changing the Tier of an Architecture

Systems may only be moved from one tier to another by approval of the FreeBSD Core Team, which shall make that decision in collaboration with the Security Officer, Release Engineering, and toolchain maintenance teams.

# 18. Ports Specific FAQ

## 18.1. Adding a New Port

Q: How do I add a new port?

A: First, please read the section about repository copies.

The easiest way to add a new port is the `addport` script located in the `ports/Tools/scripts` directory. It adds a port from the directory specified, determining the category automatically from the port Makefile. It also adds an entry to the port's category Makefile. It was written by Michael Haro <[mharo@FreeBSD.org](mailto:mharo@FreeBSD.org)>, Will Andrews <[will@FreeBSD.org](mailto:will@FreeBSD.org)>, and Renato Botelho <[garga@FreeBSD.org](mailto:garga@FreeBSD.org)>. When sending questions about this script to the [FreeBSD ports mailing list](#), please also CC Chris Rees <[crees@FreeBSD.org](mailto:crees@FreeBSD.org)>, the current maintainer.

Q: Any other things I need to know when I add a new port?

A: Check the port, preferably to make sure it compiles and packages correctly. This is the recommended sequence:

```
# make install
# make package
# make deinstall
# pkg_add package you built above
# make deinstall
# make reinstall
# make package
```

The [Porters Handbook](#) contains more detailed instructions.

Use [portlint\(1\)](#) to check the syntax of the port. You do not necessarily have to eliminate all warnings but make sure you have fixed the simple ones.

If the port came from a submitter who has not contributed to the Project before, add that person's name to the [Additional Contributors](#) section of the FreeBSD Contributors List.

Close the PR if the port came in as a PR. To close a PR, change the state to Issue Resolved and the resolution as Fixed.

## 18.2. Removing an Existing Port

Q: How do I remove an existing port?

A: First, please read the section about repository copies. Before you remove the port, you have to verify there are no other ports depending on it.

- Make sure there is no dependency on the port in the ports collection:
  - The port's PKGNAME should appear in exactly one line in a recent INDEX file.
  - No other ports should contain any reference to the port's directory or PKGNAME in their Makefiles
- Then, remove the port:
  1. Remove the port's files and directory with `svn remove`.
  2. Remove the SUBDIR listing of the port in the parent directory Makefile.
  3. Add an entry to `ports/MOVED` .
  4. Remove the port from `ports/LEGAL` if it is there.

Alternatively, you can use the `rmport` script, from `ports/Tools/scripts` . This script was written by Vasil Dimov <[vd@FreeBSD.org](mailto:vd@FreeBSD.org)>. When sending questions about this script to the [FreeBSD ports mailing list](#), please also CC Chris Rees <[crees@FreeBSD.org](mailto:crees@FreeBSD.org)>, the current maintainer.

## 18.3. Re-adding a Deleted Port

Q: How do I re-add a deleted port?

A: This is essentially the reverse of deleting a port.

1. Figure out when the port was removed. Use this [list](#) and then copy the last living revision of the port:

```
% cd /usr/ports/ category
% svn cp 'svn+ssh://svn.freebsd.org/ports/
head/ category /portname /@{YYYY-MM-DD }' portname
```

Pick a date that is before the removal but after the last true commit.

2. Perform whatever changes are necessary to make the port work again. If it was deleted because the distfiles are no longer available you will need to volunteer to host them yourself, or find someone else to do so.
3. `svn add` or `svn remove` any appropriate files.
4. Restore the SUBDIR listing of the port in the parent directory Makefile, and delete the entry from `ports/MOVED` .
5. If the port had an entry in `ports/LEGAL` , restore it.
6. `svn commit` these changes, preferably in one step.



### Tip

The `addport` script mentioned in [Q & A 18.1](#), “[Adding a New Port](#)” now detects when the port to add has previously existed, and attempts to handle all except the `ports/LEGAL` step automatically.

## 18.4. Repository Copies

Q: When do we need a repository copy?

A: When you want to add a port that is related to any port that is already in the tree in a separate directory, you have to do a repository copy. Here *related* means it is a different version or a slightly modified version. Examples are `print/ghostscript*` (different versions) and `x11-wm/windowmaker*` (English-only and internationalized version).

Another example is when a port is moved from one subdirectory to another, or when you want to change the name of a directory because the author(s) renamed their software even though it is a descendant of a port already in a tree.

Q: What do I need to do?

A: With Subversion, a repo copy can be done by any committer:

- Doing a repo copy:
  1. Verify that the target directory does not exist.

2. Use `svn up` to make certain the original files, directories, and checkout information is current.
  3. Use `svn move` or `svn copy` to do the repo copy.
  4. Upgrade the copied port to the new version. Remember to add or change the `PKGNAMEPREFIX` or `PKGNAME_SUFFIX` so there are no duplicate ports with the same name. In some rare cases it may be necessary to change the `PORT-NAME` instead of adding `PKGNAMEPREFIX` or `PKGNAME_SUFFIX`, but this should only be done when it is really needed — e.g., using an existing port as the base for a very similar program with a different name, or upgrading a port to a new upstream version which actually changes the distribution name, like the transition from `textproc/libxml` to `textproc/libxml2`. In most cases, adding or changing `PKGNAMEPREFIX` or `PKGNAME_SUFFIX` should suffice.
  5. Add the new subdirectory to the `SUBDIR` listing in the parent directory `Makefile`. You can run `make checksubdirs` in the parent directory to check this.
  6. If the port changed categories, modify the `CATEGORIES` line of the port's `Makefile` accordingly
  7. Add an entry to `ports/MOVED`, if you remove the original port.
  8. Commit all changes on one commit.
- When removing a port:
    1. Perform a thorough check of the ports collection for any dependencies on the old port location/name, and update them. Running `grep` on `INDEX` is not enough because some ports have dependencies enabled by compile-time options. A full `grep -r` of the ports collection is recommended.
    2. Remove the old port and the old `SUBDIR` entry.
    3. Add an entry to `ports/MOVED`.
  - After repo moves (“rename” operations where a port is copied and the old location is removed):
    - Follow the same steps that are outlined in the previous two entries, to activate the new location of the port and remove the old one.

## 18.5. Ports Freeze

Q: What is a “ports freeze”?



- A: Before a release, it is necessary to restrict commits to the ports tree for a short period of time while the packages and the release itself are being built. This is to ensure consistency among the various parts of the release, and is called the “ports freeze”.

For more information on the background and policies surrounding a ports freeze, see the [Portmgr Quality Assurance page](#).

- Q: What is a “ports slush” or “feature freeze”?

- A: During a release cycle the ports tree may be in a “slush” state instead of in a hard freeze. The goal during a slush is to reach a stable ports tree to avoid rebuilding large sets of packages for the release and to tag the tree. During this time “sweeping changes” are prohibited unless specifically permitted by portmgr. Complete details about what qualifies as a sweeping change can be found on the [Portmgr Implementation page](#).

The benefit of a slush as opposed to a complete freeze is that it allows maintainers to continue adding new ports, making routine version updates, and bug fixes to most existing ports, as long as the number of affected ports is minimal. For example, updating the shared library version on a port that many other ports depend on.

- Q: How long is a ports freeze or slush?

- A: A freeze only lasts long enough to tag the tree. A slush usually lasts a week or two, but may last longer.

- Q: What does it mean to me?

- A: During a ports freeze, you are not allowed to commit anything to the tree without explicit approval from the Ports Management Team. “Explicit approval” here means that you send a patch to the Ports Management Team for review and get a reply saying, “Go ahead and commit it.”

Not everything is allowed to be committed during a freeze. Please see the [Portmgr Quality Assurance page](#) for more information.

Note that you do not have implicit permission to fix a port during the freeze just because it is broken.

During a ports slush, you are still allowed to commit but you must exercise more caution in what you commit. Furthermore a special note (typically “Feature Safe: yes”) must be added to the commit message.

- Q: How do I know when the ports slush starts?

- A: The Ports Management Team will send out warning messages to the [FreeBSD ports mailing list](#) and FreeBSD committer's mailing list announcing the start of the im-

pending release, usually two or three weeks in advance. The exact starting time will not be determined until a few days before the actual release. This is because the ports slush has to be synchronized with the release, and it is usually not known until then when exactly the release will be rolled.

When the slush starts, there will be another announcement to the [FreeBSD ports mailing list](#) and FreeBSD committer's mailing list, of course.

Q: How do I know when the freeze or slush ends?

A: A few hours after the release, the Ports Management Team will send out a mail to the [FreeBSD ports mailing list](#) and FreeBSD committer's mailing list announcing the end of the ports freeze or slush. Note that the release being cut does not automatically indicate the end of the freeze. We have to make sure there will be no last minute snafus that result in an immediate re-rolling of the release.

## 18.6. Creating a New Category

Q: What is the procedure for creating a new category?

A: Please see [Proposing a New Category](#) in the Porter's Handbook. Once that procedure has been followed and the PR has been assigned to Ports Management Team <[portmgr@FreeBSD.org](mailto:portmgr@FreeBSD.org)>, it is their decision whether or not to approve it. If they do, it is their responsibility to do the following:

1. Perform any needed moves. (This only applies to physical categories.)
2. Update the `VALID_CATEGORIES` definition in `ports/Mk/bsd.port.mk`.
3. Assign the PR back to you.

Q: What do I need to do to implement a new physical category?

A: 1. Upgrade each moved port's `Makefile`. Do not connect the new category to the build yet.

To do this, you will need to:

1. Change the port's `CATEGORIES` (this was the point of the exercise, remember?) The new category should be listed *first*. This will help to ensure that the `PKGORIGIN` is correct.
2. Run a `make describe`. Since the top-level `make index` that you will be running in a few steps is an iteration of `make describe` over the entire ports hierarchy, catching any errors here will save you having to re-run that step later on.

3. If you want to be really thorough, now might be a good time to run [portlint\(1\)](#).
2. Check that the `PKGORIGIN` s are correct. The ports system uses each port's `CATEGORIES` entry to create its `PKGORIGIN` , which is used to connect installed packages to the port directory they were built from. If this entry is wrong, common port tools like [pkg\\_version\(1\)](#) and [portupgrade\(1\)](#) fail.

To do this, use the `chkorigin.sh` tool, as follows: `env PORTSDIR= /path/to/ports sh -e /path/to/ports /Tools/scripts/chkorigin.sh` . This will check *every* port in the ports tree, even those not connected to the build, so you can run it directly after the move operation. Hint: do not forget to look at the `PKGORIGIN` s of any slave ports of the ports you just moved!

3. On your own local system, test the proposed changes: first, comment out the `SUBDIR` entries in the old ports' categories' Makefiles; then enable building the new category in `ports/Makefile` . Run `make checksubdirs` in the affected category directories to check the `SUBDIR` entries. Next, in the `ports/` directory, run `make index` . This can take over 40 minutes on even modern systems; however, it is a necessary step to prevent problems for other people.
4. Once this is done, you can commit the updated `ports/Makefile` to connect the new category to the build and also commit the `Makefile` changes for the old category or categories.
5. Add appropriate entries to `ports/MOVED` .
6. Update the documentation by modifying the following:
  - the [list of categories](#) in the Porter's Handbook
  - `www/en/ports/categories` . Note that these are now displayed by subgroups, as specified in `www/en/ports/categories.descriptions` .

(Note: these are in the docs, not the ports, repository). If you are not a docs committer, you will need to submit a PR for this.

7. Only once all the above have been done, and no one is any longer reporting problems with the new ports, should the old ports be deleted from their previous locations in the repository.

It is not necessary to manually update the [ports web pages](#) to reflect the new category. This is now done automatically via your change to `www/en/ports/categories` and the daily automated rebuild of `INDEX`.

Q: What do I need to do to implement a new virtual category?

A: This is much simpler than a physical category. You only need to modify the following:

- the [list of categories](#) in the Porter's Handbook
- [www/en/ports/categories](http://www/en/ports/categories)

## 18.7. Miscellaneous Questions

Q: How do I know if my port is building correctly or not?

A: First, go check <http://pointyhat.FreeBSD.org/errorlogs/>. There you will find error logs from the latest package building runs on all supported platforms for the most recent branches.

However, just because the port does not show up there does not mean it is building correctly. (One of the dependencies may have failed, for instance.) The relevant directories are available on pointyhat under `/a/portbuild/<arch>/<major_version>` so feel free to dig around. Each architecture and version has the following subdirectories:

<code>errors</code>	error logs from latest <code>&lt;major_version&gt;</code> run on <code>&lt;arch&gt;</code>
<code>&lt;arch&gt;</code>	
<code>logs</code>	all logs from latest <code>&lt;major_version&gt;</code> run on <code>&lt;arch&gt;</code>
<code>packages</code>	packages from latest <code>&lt;major_version&gt;</code> run on <code>&lt;arch&gt;</code>
<code>bak/errors</code>	error logs from last complete <code>&lt;major_version&gt;</code> run on <code>&lt;arch&gt;</code>
<code>on &lt;arch&gt;</code>	
<code>bak/logs</code>	all logs from last complete <code>&lt;major_version&gt;</code> run on <code>&lt;arch&gt;</code>
<code>on &lt;arch&gt;</code>	
<code>bak/packages</code>	packages from last complete <code>&lt;major_version&gt;</code> run on <code>&lt;arch&gt;</code>
<code>on &lt;arch&gt;</code>	

Basically, if the port shows up in packages, or it is in logs but not in errors, it built fine. (The errors directories are what you get from the web page.)

Q: I added a new port. Do I need to add it to the INDEX?

A: No. The file can either be generated by running `make index`, or a pre-generated version can be downloaded with `make fetchindex`.

Q: Are there any other files I am not allowed to touch?

A: Any file directly under `ports/`, or any file under a subdirectory that starts with an uppercase letter (`Mk/`, `Tools/`, etc.). In particular, the Ports Management Team is very protective of `ports/Mk/bsd.port*.mk` so do not commit changes to those files unless you want to face their wrath.

Q: What is the proper procedure for updating the checksum for a port's distfile when the file changes without a version change?

- A: When the checksum for a port's distfile is updated due to the author updating the file without changing the port's revision, the commit message should include a summary of the relevant diffs between the original and new distfile to ensure that the distfile has not been corrupted or maliciously altered. If the current version of the port has been in the ports tree for a while, a copy of the old distfile will usually be available on the ftp servers; otherwise the author or maintainer should be contacted to find out why the distfile has changed.

## 19. Issues Specific to Developers Who Are Not Committers

A few people who have access to the FreeBSD machines do not have commit bits. Almost all of this document will apply to these developers as well (except things specific to commits and the mailing list memberships that go with them). In particular, we recommend that you read:

- [Administrative Details](#)
- [Conventions](#)



### Note

You should get your mentor to add you to the “Additional Contributors” (`doc/en_US.IS08859-1/articles/contributors/contrib.additional.xml`), if you are not already listed there.

- [Developer Relations](#)
- [SSH Quick-Start Guide](#)
- [The FreeBSD Committers' Big List of Rules](#)

## 20. Information About Google Analytics

As of December 12, 2012, Google Analytics was enabled on the FreeBSD Project website to collect anonymized usage statistics regarding usage of the site. The information collected is valuable to the FreeBSD Documentation Project, in order to identify various problems on the FreeBSD website.

## 20.1. Google Analytics General Policy

The FreeBSD Project takes visitor privacy very seriously. As such, the FreeBSD Project website honors the “Do Not Track” header *before* fetching the tracking code from Google. For more information, please see the [FreeBSD Privacy Policy](#).

Google Analytics access is *not* arbitrarily allowed — access must be requested, voted on by the Documentation Engineering Team <[doceng@FreeBSD.org](mailto:doceng@FreeBSD.org)>, and explicitly granted.

Requests for Google Analytics data must include a specific purpose. For example, a valid reason for requesting access would be “to see the most frequently used web browsers when viewing FreeBSD web pages to ensure page rendering speeds are acceptable.”

Conversely, “to see what web browsers are most frequently used” (without stating *why*) would be rejected.

All requests must include the timeframe for which the data would be required. For example, it must be explicitly stated if the requested data would be needed for a timeframe covering a span of 3 weeks, or if the request would be one-time only.

Any request for Google Analytics data without a clear, reasonable reason beneficial to the FreeBSD Project will be rejected.

## 20.2. Data Available Through Google Analytics

A few examples of the types of Google Analytics data available include:

- Commonly used web browsers
- Page load times
- Site access by language

## 21. Perks of the Job

Unfortunately, there are not many perks involved with being a committer. Recognition as a competent software engineer is probably the only thing that will be of benefit in the long run. However, there are at least some perks:

### Free 4-CD and DVD Sets

FreeBSD committers can get a free 4-CD or DVD set at conferences from [FreeBSD Mall, Inc.](#). The sets are no longer available as a subscription due to the high shipment costs to countries outside the USA.

### Freenode IRC Cloaks

FreeBSD developers may request a cloaked hostmask for their account on the Freenode IRC network in the form of `freebsd/developer/freefall name` or `freeb-`

sd/developer/*NickServ name*. To request a cloak, send an email to Eitan Adler <[eadler@FreeBSD.org](mailto:eadler@FreeBSD.org)> with your requested hostmask and NickServ account name.

## 22. Miscellaneous Questions

- Q: Why are trivial or cosmetic changes to files on a vendor branch a bad idea?
- A:
- From now on, every new vendor release of that file will need to have patches merged in by hand.
  - From now on, every new vendor release of that file will need to have patches *verified* by hand.
- Q: How do I add a new file to a branch?
- A: To add a file onto a branch, simply checkout or update to the branch you want to add to and then add the file using the add operation as you normally would. This works fine for the doc and ports trees. The src tree uses SVN and requires more care because of the mergeinfo properties. See section 1.4.6 of the [Subversion Primer](#) for details. Refer to [SubversionPrimer/Merging](#) for details on how to perform an MFC.
- Q: How do I access [people.FreeBSD.org](http://people.FreeBSD.org) to put up personal or project information?
- A: [people.FreeBSD.org](http://people.FreeBSD.org) is the same as [freefall.FreeBSD.org](http://freefall.FreeBSD.org). Just create a `public_html` directory. Anything you place in that directory will automatically be visible under <http://people.FreeBSD.org/>.
- Q: Where are the mailing list archives stored?
- A: The mailing lists are archived under `/g/mail` which will show up as `/hub/g/mail` with [pwd\(1\)](#). This location is accessible from any machine on the FreeBSD cluster.
- Q: I would like to mentor a new committer. What process do I need to follow?
- A: See the [New Account Creation Procedure](#) document on the internal pages.

