

Pacemaker Development

Working with the Pacemaker Code Base



Andrew Beekhof

Ken Gaillot

Pacemaker Development

Working with the Pacemaker Code Base

Edition 2

Author
Author

Andrew Beekhof
Ken Gaillot

andrew@beekhof.net
kgaillet@redhat.com

Copyright © 2016-2018 Andrew Beekhof.

The text of and illustrations in this document are licensed under version 4.0 or later of the Creative Commons Attribution-ShareAlike International Public License ("CC-BY-SA")¹.

In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

In addition to the requirements of this license, the following activities are looked upon favorably:

1. If you are distributing Open Publication works on hardcopy or CD-ROM, you provide email notification to the authors of your intent to redistribute at least thirty days before your manuscript or media freeze, to give the authors time to provide updated documents. This notification should describe modifications, if any, made to the document.
2. All substantive modifications (including deletions) be either clearly marked up in the document or else described in an attachment to the document.
3. Finally, while it is not mandatory under this license, it is considered good form to offer a free copy of any hardcopy or CD-ROM expression of the author(s) work.

This document has guidelines and tips for developers interested in editing Pacemaker source code and submitting changes for inclusion in the project.

¹ An explanation of CC-BY-SA is available at <https://creativecommons.org/licenses/by-sa/4.0/>

Table of Contents

1. Frequently Asked Questions	1
Frequently Asked Questions	1
2. C Coding Guidelines	5
2.1. C Boilerplate	5
2.2. Formatting	5
2.2.1. Whitespace	5
2.2.2. Line Length	6
2.2.3. Pointers	6
2.2.4. Functions	6
2.2.5. Control Statements (if, else, while, for, switch)	6
2.2.6. Operators	7
2.3. Naming Conventions	7
2.4. vim Settings	7
3. Python Coding Guidelines	9
3.1. Python Boilerplate	9
3.2. Python Compatibility	10
3.2.1. Python Future Imports	10
3.2.2. Other Python Compatibility Requirements	10
3.2.3. Python Usages to Avoid	11
3.3. Formatting Python Code	11
A. Revision History	13
Index	15

Frequently Asked Questions

Q: Who is this document intended for?

A: Anyone who wishes to read and/or edit the Pacemaker source code. Casual contributors should feel free to read just this FAQ, and consult other sections as needed.

Q: Where is the source code for Pacemaker?

A: The [source code for Pacemaker](#)¹ is kept on [GitHub](#)², as are all software projects under the [ClusterLabs](#)³ umbrella. Pacemaker uses [Git](#)⁴ for source code management. If you are a Git newbie, the [gittutorial\(7\) man page](#)⁵ is an excellent starting point. If you're familiar with using Git from the command line, you can create a local copy of the Pacemaker source code with: **git clone https://github.com/ClusterLabs/pacemaker.git pacemaker**

Q: What are the different Git branches and repositories used for?

A:

- The [master branch](#)⁶ is the primary branch used for development.
 - The [1.1 branch](#)⁷ contains the latest official release, and normally does not receive any changes. During the release cycle, it will contain release candidates for the next official release, and will receive only bug fixes.
 - The [1.0 repository](#)⁸ is a frozen snapshot of the 1.0 release series, and is no longer developed.
 - Messages will be posted to the developers@clusterlabs.org⁹ mailing list during the release cycle, with instructions about which branches to use when submitting requests.
-

Q: How do I build from the source code?

A: See [INSTALL.md](#)¹⁰ in the main checkout directory.

Q: What coding style should I follow?

A: You'll be mostly fine if you simply follow the example of existing code. When unsure, see the relevant section of this document for language-specific recommendations. Pacemaker has grown and evolved organically over many years, so you will see much code that doesn't conform to the current guidelines. We discourage making changes solely to bring code into conformance, as any change requires developer time for review and opens the possibility of adding bugs. However, new code should follow the guidelines, and it is fine to bring lines of older code into conformance when modifying that code for other reasons.

¹ <https://github.com/ClusterLabs/pacemaker>

² <https://github.com/>

³ <https://github.com/ClusterLabs>

⁴ <https://git-scm.com/>

⁵ <http://schacon.github.io/git/gittutorial.html>

⁶ <https://github.com/ClusterLabs/pacemaker/tree/master>

⁷ <https://github.com/ClusterLabs/pacemaker/tree/1.1>

⁸ <https://github.com/ClusterLabs/pacemaker-1.0>

⁹ <http://clusterlabs.org/mailman/listinfo/developers>

¹⁰ <https://github.com/ClusterLabs/pacemaker/blob/master/INSTALL.md>

Chapter 1. Frequently Asked Questions

Q: How should I format my Git commit messages?

A: See existing examples in the git log. The first line should look like **change-type: affected-code: explanation** where **change-type** can be **Fix** or **Bug** for most bug fixes, **Feature** for new features, **Log** for changes to log messages or handling, **Doc** for changes to documentation or comments, or **Test** for changes in CTS and regression tests. You will sometimes see **Low**, **Med** (or **Mid**) and **High** used instead for bug fixes, to indicate the severity. The important thing is that only commits with **Feature**, **Fix**, **Bug**, or **High** will automatically be included in the change log for the next release. The **affected-code** is the name of the component(s) being changed, for example, **pacemaker-controld** or **libcrmcommon** (it's more free-form, so don't sweat getting it exact). The **explanation** briefly describes the change. The git project recommends the entire summary line stay under 50 characters, but more is fine if needed for clarity. Except for the most simple and obvious of changes, the summary should be followed by a blank line and then a longer explanation of *why* the change was made.

Q: How can I test my changes?

A: Most importantly, Pacemaker has regression tests for most major components; these will automatically be run for any pull requests submitted through GitHub. Additionally, Pacemaker's Cluster Test Suite (CTS) can be used to set up a test cluster and run a wide variety of complex tests. This document will have more detail on testing in the future.

Q: What is Pacemaker's license?

A: Except where noted otherwise in the file itself, the source code for all Pacemaker programs is licensed under version 2 or later of the GNU General Public License ([GPLv2+](https://www.gnu.org/licenses/gpl-2.0.html)¹¹), its headers and libraries under version 2.1 or later of the less restrictive GNU Lesser General Public License ([LGPLv2.1+](https://www.gnu.org/licenses/lgpl-2.1.html)¹²), its documentation under version 4.0 or later of the Creative Commons Attribution-ShareAlike International Public License ([CC-BY-SA](https://creativecommons.org/licenses/by-sa/4.0/legalcode)¹³), and its init scripts under the [Revised BSD](https://opensource.org/licenses/BSD-3-Clause)¹⁴ license. If you find any deviations from this policy, or wish to inquire about alternate licensing arrangements, please e-mail andrew@beekhof.net¹⁵. Licensing issues are also discussed on the [ClusterLabs wiki](http://clusterlabs.org/wiki)¹⁶.

Q: How can I contribute my changes to the project?

A: Contributions of bug fixes or new features are very much appreciated! Patches can be submitted as [pull requests](https://help.github.com/articles/using-pull-requests/)¹⁷ via GitHub (the preferred method, due to its excellent [features](https://github.com/features/)¹⁸), or e-mailed to the [developers@clusterlabs.org](http://clusterlabs.org/mailman/listinfo/developers)¹⁹ mailing list as an attachment in a format Git can import.

Q: What if I still have questions?

¹¹ <https://www.gnu.org/licenses/gpl-2.0.html>

¹² <https://www.gnu.org/licenses/lgpl-2.1.html>

¹³ <https://creativecommons.org/licenses/by-sa/4.0/legalcode>

¹⁴ <https://opensource.org/licenses/BSD-3-Clause>

¹⁵ <mailto:andrew@beekhof.net>

¹⁶ <http://clusterlabs.org/wiki/License>

¹⁷ <https://help.github.com/articles/using-pull-requests/>

¹⁸ <https://github.com/features/>

¹⁹ <http://clusterlabs.org/mailman/listinfo/developers>

A: Ask on the developers@clusterlabs.org²⁰ mailing list for development-related questions, or on the users@clusterlabs.org²¹ mailing list for general questions about using Pacemaker. Developers often also hang out on [freenode's](http://freenode.net)²² #clusterlabs IRC channel.

²⁰ <http://clusterlabs.org/mailman/listinfo/developers>

²¹ <http://clusterlabs.org/mailman/listinfo/users>

²² <http://freenode.net/>

C Coding Guidelines

Table of Contents

2.1. C Boilerplate	5
2.2. Formatting	5
2.2.1. Whitespace	5
2.2.2. Line Length	6
2.2.3. Pointers	6
2.2.4. Functions	6
2.2.5. Control Statements (if, else, while, for, switch)	6
2.2.6. Operators	7
2.3. Naming Conventions	7
2.4. vim Settings	7

2.1. C Boilerplate

Every C file should start like this:

```
/*
 * Copyright <YYYY[-YYYY]> Andrew Beekhof <andrew@beekhof.net>
 *
 * This source code is licensed under <LICENSE> WITHOUT ANY WARRANTY.
 */
```

<YYYY> is the year the code was *originally* created.¹ If the code is modified in later years, add -YYYY with the most recent year of modification.

<LICENSE> should follow the policy set forth in the [COPYING](#)² file, generally one of "GNU General Public License version 2 or later (GPLv2+)" or "GNU Lesser General Public License version 2.1 or later (LGPLv2.1+)".

2.2. Formatting

2.2.1. Whitespace

- Indentation must be 4 spaces, no tabs.
- Do not leave trailing whitespace.

¹ See the U.S. Copyright Office's "[Compendium of U.S. Copyright Office Practices](https://www.copyright.gov/comp3/)" [https://www.copyright.gov/comp3/], particularly "Chapter 2200: Notice of Copyright", sections 2205.1(A) and 2205.1(F), or "[Updating Copyright Notices](https://techwhirl.com/updates-copyright-notices/)" [https://techwhirl.com/updates-copyright-notices/] for a more readable summary.

² <https://github.com/ClusterLabs/pacemaker/blob/master/COPYING>

2.2.2. Line Length

- Lines should be no longer than 80 characters unless limiting line length significantly impacts readability.

2.2.3. Pointers

- The * goes by the variable name, not the type:

```
char *foo;
```

- Use a space before the * and after the closing parenthesis in a cast:

```
char *foo = (char *) bar;
```

2.2.4. Functions

- In the function definition, put the return type on its own line, and place the opening brace by itself on a line:

```
static int  
foo(void)  
{
```

- For functions with enough arguments that they must break to the next line, align arguments with the first argument:

```
static int  
function_name(int bar, const char *a, const char *b,  
              const char *c, const char *d)  
{
```

- If a function name gets really long, start the arguments on their own line with 8 spaces of indentation:

```
static int  
really_really_long_function_name_this_is_getting_silly_now(  
    int bar, const char *a, const char *b,  
    const char *c, const char *d)  
{
```

2.2.5. Control Statements (if, else, while, for, switch)

- The keyword is followed by one space, then left parenthesis without space, condition, right parenthesis, space, opening bracket on the same line. **else** and **else if** are on the same line with the ending brace and opening brace, separated by a space:

```
if (condition1) {  
    statement1;  
} else if (condition2) {  
    statement2;
```

```
} else {  
    statement3;  
}
```

- In a **switch** statement, **case** is indented one level, and the body of each **case** is indented by another level. The opening brace is on the same line as **switch**.

```
switch (expression) {  
    case 0:  
        command1;  
        break;  
    case 1:  
        command2;  
        break;  
    default:  
        command3;  
}
```

2.2.6. Operators

- Operators have spaces from both sides. Do not rely on operator precedence; use parentheses when mixing operators with different priority.
- No space is used after opening parenthesis and before closing parenthesis.

```
x = a + b - (c * d);
```

2.3. Naming Conventions

- Any exposed symbols in libraries (non-**static** function names, type names, etc.) must begin with a prefix appropriate to the library, for example, **crm_**, **pe_**, **st_**, **lrm_**.

2.4. vim Settings

Developers who use **vim** to edit source code can add the following settings to their `~/ .vimrc` file to follow Pacemaker C coding guidelines:

```
" follow Pacemaker coding guidelines when editing C source code files  
filetype plugin indent on  
au FileType c    setlocal expandtab tabstop=4 softtabstop=4 shiftwidth=4 textwidth=80  
autocmd BufNewFile,BufRead *.h set filetype=c  
let c_space_errors = 1
```


Python Coding Guidelines

Table of Contents

3.1. Python Boilerplate	9
3.2. Python Compatibility	10
3.2.1. Python Future Imports	10
3.2.2. Other Python Compatibility Requirements	10
3.2.3. Python Usages to Avoid	11
3.3. Formatting Python Code	11

3.1. Python Boilerplate

If a Python file is meant to be executed (as opposed to imported), it should have a `.in` extension, and its first line should be:

```
#!@PYTHON@
```

which will be replaced with the appropriate python executable when Pacemaker is built. To make that happen, add an `AC_CONFIG_FILES()` line to `configure.ac`, and add the file name without `.in` to `.gitignore` (see existing examples).

After the above line if any, every Python file should start like this:

```
""" <BRIEF-DESCRIPTION>
"""

# Pacemaker targets compatibility with Python 2.7 and 3.2+
from __future__ import print_function, unicode_literals, absolute_import, division

__copyright__ = "Copyright <YYYY[-YYYY]> Andrew Beekhof <andrew@beekhof.net>"
__license__ = "<LICENSE> WITHOUT ANY WARRANTY"
```

If the file is meant to be directly executed, the first line (**<SHEBANG>**) should be `#!/usr/bin/python`. If it is meant to be imported, omit this line.

<BRIEF-DESCRIPTION> is obviously a brief description of the file's purpose. The string may contain any other information typically used in a Python file *docstring*¹.

The `import` statement is discussed further in [Section 3.2.1, "Python Future Imports"](#).

<YYYY> is the year the code was *originally* created.² If the code is modified in later years, add `-YYYY` with the most recent year of modification.

¹ <https://www.python.org/dev/peps/pep-0257/>

² See the U.S. Copyright Office's ["Compendium of U.S. Copyright Office Practices"](https://www.copyright.gov/comp3/) [https://www.copyright.gov/comp3/], particularly "Chapter 2200: Notice of Copyright", sections 2205.1(A) and 2205.1(F), or ["Updating Copyright Notices"](https://techwhirl.com/updates/copyright-notices/) [https://techwhirl.com/updates/copyright-notices/] for a more readable summary.

<LICENSE> should follow the policy set forth in the [COPYING](#)³ file, generally one of "GNU General Public License version 2 or later (GPLv2+)" or "GNU Lesser General Public License version 2.1 or later (LGPLv2.1+)".

3.2. Python Compatibility

Pacemaker targets compatibility with Python 2.7, and Python 3.2 and later. These versions have added features to be more compatible with each other, allowing us to support both the 2 and 3 series with the same code. It is a good idea to test any changes with both Python 2 and 3.

3.2.1. Python Future Imports

The future imports used in [Section 3.1, "Python Boilerplate"](#) mean:

- All print statements must use parentheses, and printing without a newline is accomplished with the `end=' '` parameter rather than a trailing comma.
- All string literals will be treated as Unicode (the `u` prefix is unnecessary, and must not be used, because it is not available in Python 3.2).
- Local modules must be imported using `from . import` (rather than just `import`). To import one item from a local module, use `from .modulename import` (rather than `from modulename import`).
- Division using `/` will always return a floating-point result (use `//` if you want the integer floor instead).

3.2.2. Other Python Compatibility Requirements

- When specifying an exception variable, always use `as` instead of a comma (e.g. `except Exception as e` or `except (TypeError, IOError) as e`). Use `e.args` to access the error arguments (instead of iterating over or subscripting `e`).
- Use `in` (not `has_key()`) to determine if a dictionary has a particular key.
- Always use the I/O functions from the `io` module rather than the native I/O functions (e.g. `io.open()` rather than `open()`).
- When opening a file, always use the `t` (text) or `b` (binary) mode flag.
- When creating classes, always specify a parent class to ensure that it is a "new-style" class (e.g. `class Foo(object):` rather than `class Foo:`).
- Be aware of the bytes type added in Python 3. Many places where strings are used in Python 2 use bytes or bytearrays in Python 3 (for example, the pipes used with `subprocess.Popen()`). Code should handle both possibilities.
- Be aware that the `items()`, `keys()`, and `values()` methods of dictionaries return lists in Python 2 and views in Python 3. In many case, no special handling is required, but if the code needs to use list methods on the result, cast the result to list first.
- Do not raise or catch strings as exceptions (e.g. `raise "Bad thing"`).

³ <https://github.com/ClusterLabs/pacemaker/blob/master/COPYING>

- Do not use the `cmp` parameter of sorting functions (use `key` instead, if needed) or the `__cmp__()` method of classes (implement rich comparison methods such as `__lt__()` instead, if needed).
- Do not use the `buffer` type.
- Do not use features not available in all targeted Python versions. Common examples include:
 - The `html`, `ipaddress`, and `UserDict` modules
 - The `subprocess.run()` function
 - The `subprocess.DEVNULL` constant
 - `subprocess` module-specific exceptions

3.2.3. Python Usages to Avoid

Avoid the following if possible, otherwise research the compatibility issues involved (hacky workarounds are often available):

- long integers
- octal integer literals
- mixed binary and string data in one data file or variable
- metaclasses
- `locale.strcoll` and `locale.strxfrm`
- the `configparser` and `ConfigParser` modules
- importing compatibility modules such as `six` (so we don't have to add them to Pacemaker's dependencies)

3.3. Formatting Python Code

- Indentation must be 4 spaces, no tabs.
- Do not leave trailing whitespace.
- Lines should be no longer than 80 characters unless limiting line length significantly impacts readability. For Python, this limitation is flexible since breaking a line often impacts readability, but definitely keep it under 120 characters.
- Where not conflicting with this style guide, it is recommended (but not required) to follow [PEP 8](#)⁴.
- It is recommended (but not required) to format Python code such that `pylint --disable=line-too-long,too-many-lines,too-many-instance-attributes,too-many-arguments,too-many-statements` produces minimal complaints (even better if you don't need to disable all those checks).

⁴ <https://www.python.org/dev/peps/pep-0008/>

Appendix A. Revision History

Revision 1-0 **Tue Jul 26 2016** **Ken Gaillot** kgaillot@redhat.com

Convert coding guidelines and developer FAQ to Publican document

Revision 1-1 **Mon Aug 29 2016** **Ken Gaillot** kgaillot@redhat.com

Add Python coding guidelines, and more about licensing

Revision 2-0 **Fri Jan 12 2018** **Ken Gaillot** kgaillot@redhat.com

Drop support for Python 2.6

Index

Symbols

2, 10
3, 10

B

boilerplate, 5, 9
branches, 1

C

C
boilerplate, 5
functions, 6
naming, 7
operators, 7
pointers, 6
whitespace, 5
C boilerplate, 5
commit messages, 2

D

downloads, 1

F

formatting, 11
functions, 6

G

git
commit messages, 2
GitHub, 1
GitHub, 1

L

licensing, 2
C boilerplate, 5
Python boilerplate, 9

M

mailing lists, 3

N

naming, 7

O

operators, 7

P

pointers, 6
Python

2, 10
3, 10
boilerplate, 9
formatting, 11
versions, 10
Python boilerplate, 9

S

source code, 1

V

versions, 10
vim, 7

W

whitespace, 5
