

# Table of Contents

<a href="#">Introduction</a>	1.1
<a href="#">Legal Notice</a>	1.2
<a href="#">Working with the Code</a>	1.3
<a href="#">IDE Integration</a>	1.4
<a href="#">Building</a>	1.5
<a href="#">Tests</a>	1.6
<a href="#">Code coverage report</a>	1.7
<a href="#">Code Formatting</a>	1.8
<a href="#">Validating releases</a>	1.9
<a href="#">Notes for Maintainers</a>	1.10
<a href="#">History</a>	1.11



## Apache ActiveMQ Artemis Hacking Guide

This hacking guide outlines how developers can get involved in contributing to the Apache ActiveMQ Artemis project.

## Legal Notice

Licensed to the Apache Software Foundation (ASF) under one or more contributor license agreements. See the NOTICE file distributed with this work for additional information regarding copyright ownership. The ASF licenses this file to You under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

## Working with the Code

While the canonical Apache ActiveMQ Artemis git repository is hosted on Apache hardware at <https://gitbox.apache.org/repos/asf/activemq-artemis.git> contributors are encouraged (but not required) to use a mirror on GitHub for collaboration and pull-request review functionality. Follow the steps below to get set up with GitHub, etc.

If you do not wish to use GitHub for whatever reason you can follow the overall process outlined in the "Typical development cycle" section below but instead attach a [patch file](#) to the related JIRA or an email to the [dev list](#).

### Initial Steps

1. Create a GitHub account if you don't have one already

<https://github.com>

2. Fork the apache-artemis repository into your account

<https://github.com/apache/activemq-artemis>

3. Clone your newly forked copy onto your local workspace:

```
$ git clone git@github.com:<your-user-name>/activemq-artemis.git
Cloning into 'activemq-artemis'...
remote: Counting objects: 63800, done.
remote: Compressing objects: 100% (722/722), done.
remote: Total 63800 (delta 149), reused 0 (delta 0), pack-reused 62748
Receiving objects: 100% (63800/63800), 18.28 MiB | 3.16 MiB/s, done.
Resolving deltas: 100% (28800/28800), done.
Checking connectivity... done.

$ cd activemq-artemis
```

4. Add a remote reference to `upstream` for pulling future updates

```
$ git remote add upstream https://github.com/apache/activemq-artemis
```

5. Build with Maven

Typically developers will want to build using the `dev` profile which enables license and code style checks. For example:

```

$ mvn -Pdev install
...
[INFO] -----
[INFO] Reactor Summary:
[INFO]
[INFO] ActiveMQ Artemis Parent ..... SUCCESS [2.298
[INFO] ActiveMQ Artemis Commons ..... SUCCESS [1.821
[INFO] ActiveMQ Artemis Selector Implementation ..... SUCCESS [0.767
[INFO] ActiveMQ Artemis Native POM ..... SUCCESS [0.189
[INFO] ActiveMQ Artemis Journal ..... SUCCESS [0.646
[INFO] ActiveMQ Artemis Core Client ..... SUCCESS [5.969
[INFO] ActiveMQ Artemis JMS Client ..... SUCCESS [2.110
[INFO] ActiveMQ Artemis Server ..... SUCCESS [11.54
...
[INFO] ActiveMQ Artemis stress Tests ..... SUCCESS [0.332
[INFO] ActiveMQ Artemis performance Tests ..... SUCCESS [0.174
[INFO] -----
[INFO] BUILD SUCCESS
[INFO] -----

```

## Typical development cycle

1. Identify a task (e.g. a bug to fix or feature to implement)

<https://issues.apache.org/jira/browse/ARTEMIS>

2. Create a topic branch in your local git repo to do your work

```
$ git checkout -b my_cool_feature
```

3. Make the changes and commit one or more times

```
$ git commit
```

When you commit your changes you will need to supply a commit message. We follow the 50/72 git commit message format. An ActiveMQ Artemis commit message should be formatted in the following manner:

- i. Add the first line with the summary, using maximum 50 characters. Start the summary with the jira key (ARTEMIS-XXX) followed by a brief description of the change. Use the prefix `NO-JIRA` only for a very small insignificant change, like a typo or a small doc fix. Bug fixes, features or any code change, really should be accompanied by a jira, so they can clearly be reported in the release notes.
- ii. Insert a single blank line after the first line.
- iii. Provide a detailed description of the change in the following lines, breaking paragraphs where needed. These lines should be wrapped at 72 characters.

An example correctly formatted commit message:

```
ARTEMIS-123 Add new commit msg format to README
```

```
Adds a description of the new commit message format as well as examples
of well formatted commit messages to the README.md. This is required
to enable developers to quickly identify what the commit is intended to
do and why the commit was added.
```

4. Occasionally you'll want to push your commit(s) to GitHub for safe-keeping and/or sharing with others.

```
git push origin my_cool_feature
```

Note that `git push` references the branch you are pushing and defaults to `master`, not your working branch.

5. Discuss your planned changes (if you want feedback)

On mailing list - <http://activemq.apache.org/mailling-lists.html> On IRC - <irc://irc.freenode.org/apache-activemq> or <https://webchat.freenode.net/?channels=apache-activemq>

6. Once you're finished coding your feature/fix then rebase your branch against the latest master (applies your patches on top of master)

```
git fetch upstream
git rebase -i upstream/master
# if you have conflicts fix them and rerun rebase
# The -f, forces the push, alters history, see note below
git push -f origin my_cool_feature
```

The `rebase -i` triggers an interactive update which also allows you to combine commits, alter commit messages etc. It's a good idea to make the commit log very nice for external consumption (e.g. by squashing all related commits into a single commit). Note that rebasing and/or using `push -f` can alter history. While this is great for making a clean patch, it is unfriendly to anyone who has forked your branch. Therefore you'll want to make sure that you either work in a branch that you don't share, or if you do share it, tell them you are about to revise the branch history (and thus, they will then need to rebase on top of your branch once you push it out).

7. Get your changes merged into upstream
  - i. Send a GitHub pull request, by clicking the pull request link while in your repo's fork.
  - ii. An email will automatically be sent to the ActiveMQ developer list.
  - iii. As part of the review you may see an automated test run comment on your request.
  - iv. After review a maintainer will merge your PR into the canonical git repository at which point those changes will be synced with the GitHub mirror repository (i.e. your `master`) and your PR will be closed by the `asfgit` bot.

## Other common tasks

### 1. Pulling updates from upstream

```
$ git pull --rebase upstream master
```

(`--rebase` will automatically move your local commits, if any, on top of the latest branch you pull from; you can leave it off if you do not have any local commits).

One last option, which some prefer, is to avoid using pull altogether, and just use fetch + rebase (this is of course more typing). For example:

```
$ git fetch upstream  
$ git pull
```

### 2. Pushing pulled updates (or local commits if you aren't using topic branches) to your private GitHub repo (origin)

```
$ git push  
Counting objects: 192, done.  
Delta compression using up to 4 threads.  
Compressing objects: 100% (44/44), done.  
Writing objects: 100% (100/100), 10.67 KiB, done.  
Total 100 (delta 47), reused 100 (delta 47)  
To git@github.com:<your-user-name>/apache-artemis.git  
3382570..1fa25df master -> master
```

You might need to say `-f` to force the changes.

## Adding New Dependencies

Due to incompatibilities between some open source licenses and the Apache v2.0 license (that this project is licensed under) care must be taken when adding new dependencies to the project. The Apache Software Foundation 3rd party licensing policy has more information here: <https://www.apache.org/legal/3party.html>

To keep track of all licenses in ActiveMQ Artemis, new dependencies must be added in either the top level pom.xml or in test/pom.xml (depending on whether this is a test only dependency or if it is used in the main code base). The dependency should be added under the dependency management section with version and labelled with a comment highlighting the license for the dependency version. See existing dependencies in the main pom.xml for examples. The dependency can then be added to individual ActiveMQ Artemis modules *without* the version specified (the version is implied from the dependency management section of the top level pom). This allows ActiveMQ Artemis developers to keep track of all dependencies and licenses.

## IDE Integration

There are a few files useful for IDE integration under `./etc/ide-settings` on a checked out folder. This folder is not part of the source distribution, but it can be easily obtained:

- <https://github.com/apache/activemq-artemis/tree/master/etc/ide-settings>

## IntelliJ IDEA

### Importing the Project

The following steps show how to open the ActiveMQ Artemis maven project in IntelliJ IDEA and setup the correct maven profile to allow running of JUnit tests from within the IDE. (Steps are based on version: 2018.2)

- Starting from a clean clone
- Build the `activemq-artemis` source using the maven 'dev' profile:

```
$ mvn -Pdev install -DskipTests=true
```

- Await and ensure the outcome is BUILD SUCCESS
- Launch IntelliJ
- From the 'Welcome to IntelliJ IDEA' screen select 'Open'
- From the 'Open File or Project' screen select the `pom.xml` file in the cloned `activemq-artemis` directory
- From the 'Open Project' screen select 'Open as Project'

This should open the main IntelliJ IDEA window where you will notice some background tasks running via the bottom status bar. These background tasks should successfully import and index the project automatically.

Once the project has been imported and IntelliJ IDEA has caught up importing all the relevant dependencies, you should be able to run JUnit tests from within the IDE. Select any test class in the `tests -> integration tests` folder. Right click on the class in the project tab and click "Run". If the "Run" option is present then you're all set to go.

### Note about IBM JDK on Idea

If you are running IBM JDK it may be a little tricky to get it working.

After you add the JDK to the IDE, add also the `vm.jar` specific to your platform under that `jdk`.

```
(e.g: JAVA_HOME/jre/lib/amd64/default/jclSC180/vm.jar
```

There's a [SOF Question](#) about this that could be useful in case you are running into this issue.



## Style Templates and Inspection Settings for Idea

We have shared the style templates that are good for this project. If you want to apply them use these steps:

- File->Import Settings
- Select the file under `./artemis-cloned-folder/etc/ide-settings/idea/IDEA-style.jar`
- Select both Code Style Templates and File templates (it's the default option)
- Select OK and restart Idea

Alternatively you can copy `artemis-codestyle.xml` under your home settings at `IntelliJ IDEA15/codestyles` .

### To import inspection settings:

- File->Settings->Editor->Inspections->Manage->Import
- Select the file `./artemis-cloned-folder/etc/ide-settings/idea/artemis-inspections.xml`
- Select OK

### Issue: My JUnit tests are not runnable with in the IDE.

If the "Run " or "Run all tests" option is not present. It is likely that the default profile has not been imported properly. To (re)import the "tests" Maven profile in an existing project.

- Open the Maven Projects Tool Window: View -> Tool Windows -> Maven Projects
- Select the "profiles" drop down
- Unselect then reselect the checkbox next to "tests".
- Click on the "Reimport all maven projects" button in the top left hand corner of the window. (It looks like a circular blue arrow.
- Wait for IDEA to reload and try running a JUnit test again. The option to run should now be present.

## Eclipse

We recommend using Eclipse Kepler (4.3), due to the built-in support for Maven and Git. Note that there are still some Maven plugins used by sub-projects (e.g. documentation) which are not supported even in Eclipse Kepler (4.3).

Eclipse [m2e](#) is already included in "Eclipse IDE for Java Developers", or it can be installed from [Eclipse Kepler release repository](#).

### Git setup

It is strongly recommended to turn off the auto-updating of `.gitignore` files by the Git Team extension. Otherwise, it generates new `.gitignore` files in many directories that are not needed due to the top level `.gitignore` file. To turn it off, go

to Preferences->Team->Git->Projects and deselect the "Automatically ignore derived resources" checkbox.

## Schema setup

For proper schema validation you can add the Artemis schemas to your Eclipse XML Catalog

- Open: Window -> Preferences -> XML -> XML Catalog
- Select Add -> Workspace -> Navigate to artemis-server and select src/main/resources/schema/artemis-server.xsd -> click OK
- Repeat the above steps and add src/main/resources/schema/artemis-configuration.xsd

## Checkstyle setup

You can import the Artemis Checkstyle template into eclipse to do Checkstyle validation. As a prerequisite you need to make sure the Checkstyle plugin is installed into Eclipse which you can get from the Eclipse Marketplace. You also will need to configure Sevntu-Checkstyle. See <https://sevntu-checkstyle.github.io/sevntu.checkstyle/> for instructions. Then to configure the template:

- Open: Window -> Preferences -> Checkstyle
- Select New -> "Project Relative Configuration" in the "Type" dropdown
- Give the configuration a name and under location put "/artemis-pom/etc/checkstyle.xml" then hit ok
- You should now see your new configuration in the list of Checkstyle profiles. You can select the new configuration as the default if you want.

## Annotation Pre-Processing

ActiveMQ Artemis uses [JBoss Logging](#) and that requires source code generation from Java annotations. In order for it to 'just work' in Eclipse you need to install the *Maven Integration for Eclipse JDT Annotation Processor Toolkit m2e-apt*. See this [JBoss blog post](#) for details.

## Running tests from Eclipse

Setting up annotation pre-processing in the above section is all you need to run tests in the "unit-tests" project as that will properly add the generated logger to the source. However, one more step is needed to run tests in other projects such as "performance-tests" or "integration-tests" that have a dependency on "unit-tests". Currently m2eclipse does not properly link the generated source annotations folder from "unit-tests" which causes the logger that is generated to not be available. To simplest way to fix this is to manually add a project dependency on "unit-tests" to each of the projects where you want to run a test class from:

- Right click on the test project (i.e. integration-tests): Properties -> Java Build Path -> Projects -> Add
- Select the "unit-tests" project and click Ok

You should now be able to run tests assuming that the annotation pre-processing was set up properly in the previous step.

## M2E Connector for Javacc-Maven-Plugin

Eclipse Indigo (3.7) has out-of-the-box support for it.

As of this writing, Eclipse Kepler (4.3) still lacks support for Maven's javacc plugin. The available [m2e connector for javacc-maven-plugin](#) requires a downgrade of Maven components to be installed. manual installation instructions (as of this writing you need to use the development update site). See [this post](#) for how to do this with Eclipse Juno (4.2).

The current recommended solution for Eclipse Kepler is to mark `javacc-maven-plugin` as ignored by Eclipse, run Maven from the command line and then modify the project `activemq-core-client` adding the folder `target/generated-sources/javacc` to its build path.

## Use *Project Working Sets*

Importing all ActiveMQ Artemis subprojects will create *too many* projects in Eclipse, cluttering your *Package Explorer* and *Project Explorer* views. One way to address that is to use [Eclipse's Working Sets](#) feature. A good introduction to it can be found at a [Dzone article on Eclipse Working Sets](#).

## Building

We use Apache Maven to build the code, distribution, etc. and to manage dependencies.

We use `gitbook` to build the docs.

The minimum required Maven version is 3.0.0.

Note that there are some [compatibility issues with Maven 3.X](#) still unsolved. This is specially true for the ['site' plugin](#).

## Full Release

### Upgrading the `gitbook` version and regenerating the `npm-shrinkwrap.json` file

The full release uses `gitbook` to build a static website from the documentation. This is automatically installed using an `NPM` plugin and is controlled via a `package.json` file.

Install `NPM` using the instructions below

```
cd artemis-website
alter the `package.json` changing the version

npm cache clean; rm -rf ./node_modules/ ./node npm-shrinkwrap.json
npm install --save-dev
npm shrinkwrap --dev
```

The new `npm-shrinkwrap.json` should be written, commit it.

### `gitbook`

Artemis will automate the execution and download of `npm`. But it may be useful to install it on your system. Simply type:

```
$ npm install -g gitbook-cli
```

If you don't have `npm` installed then you would need to install it first.

### Install `npm` On Fedora

```
$ yum install npm
```

### Install `npm` On Fedora 24

This is what you should do if you are using Fedora 24+.

```
$ dnf install nodejs
```

## Install npm On Mac-OS

The easiest way would be through brew [brew](#)

You first install brew using the instructions on the [brew](#) website.

After you installed brew you can install npm by:

```
brew install npm
```

To build the full release with documentation, Javadocs, and the full web site:

```
$ mvn -Prelease package
```

To install it to your local maven repo:

```
$ mvn -Prelease install
```

## Build the distribution without docs

It is possible to build a distribution without the manuals and Javadocs. simply run

```
$ mvn package
```

## Building the docs

From either `docs/hacking-guide/en` Or `docs/user-manual/en` run `gitbook build`  
(after you've installed gitbook)

# Tests

## Running Tests

To run the unit tests:

```
$ mvn -Ptests test
```

Generating reports from unit tests:

```
$ mvn install site
```

Running tests individually

```
$ mvn -Ptests -DfailIfNoTests=false -Dtest=<test-name> test
```

where `<test-name>` is the name of the Test class without its package name

## Writing Tests

The broker is comprised of POJOs so it's simple to configure and run a broker instance and test particular functionality. Even complex test-cases involving multiple clustered brokers are relatively easy to write. Almost every test in the test-suite follows this pattern - configure broker, start broker, test functionality, stop broker.

The test-suite uses JUnit to manage test execution and life-cycle. Most tests extend `org.apache.activemq.artemis.tests.util.ActiveMQTestBase` which contains JUnit setup and tear-down methods as well as a wealth of utility functions to configure, start, manage, and stop brokers as well as perform other common tasks.

Check out `org.apache.activemq.artemis.tests.integration.SimpleTest`. It's a very simple test-case that extends

`org.apache.activemq.artemis.tests.util.ActiveMQTestBase` and uses its methods to configure a server, run a test, and then `super.tearDown()` cleans it up once the test completes. The test-case includes comments to explain everything. As the name implies, this is a simple test-case that demonstrates the most basic functionality of the test-suite. A simple test like this takes less than a second to run on modern hardware.

Although `org.apache.activemq.artemis.tests.integration.SimpleTest` is simple it could be simpler still by extending

`org.apache.activemq.artemis.tests.util.SingleServerTestBase`. This class does all the setup of a simple server automatically and provides the test-case with a `ServerLocator`, `ClientSessionFactory`, and `ClientSession` instance.

`org.apache.activemq.artemis.tests.integration.SingleServerSimpleTest` is an example based on `org.apache.activemq.artemis.tests.integration.SimpleTest` but extends `org.apache.activemq.artemis.tests.util.SingleServerTestBase` which eliminates all the setup and class variables which are provided by `SingleServerTestBase` itself.

## Keys for writing good tests

### Use log.debug

- Please use `log.debug` instead of `log.info`.

On your classes, import the following:

```
public class MyTest {
    private static final org.jboss.logging.Logger log = org.jboss.logging.Logger.getLogger(MyTest.class);

    @Test
    public void test() {
        log.debug("Log only what you need please!");
    }
}
```

- Please do not use `System.out.println()`

As a general rule, only use `System.out` if you really intend an error to be on the reporting. Debug information should be called through `log.debug`.

### Avoid leaks

An important task for any test-case is to clean up all the resources it creates when it runs. This includes the server instance itself and any resources created to connect to it (e.g. instances of `ServerLocator`, `ClientSessionFactory`, `ClientSession`, etc.). This task is typically completed in the test's `tearDown()` method. However, `ActiveMQTestBase` (and other classes which extend it) simplifies this process. As

`org.apache.activemq.artemis.tests.integration.SimpleTest` demonstrates, there are several methods you can use when creating your test which will ensure proper clean up *automatically* when the test is torn down. These include:

- All the overloaded `org.apache.activemq.artemis.tests.util.ActiveMQTestBase.createServer(..)` methods. If you choose *not* to use one of these methods to create your `ActiveMQServer` instance then use the `addServer(ActiveMQServer)` method to add the instance to the test-suite's internal resource ledger.
- Methods from `org.apache.activemq.artemis.tests.util.ActiveMQTestBase` to create a `ServerLocator` like `createInVMNonHALocator` and `createNettyNonHALocator`. If you choose *not* to use one of these methods then use `addServerLocator(ServerLocator)` to add the locator to the test-suite's internal resource ledger.

- `org.apache.activemq.artemis.tests.util.ActiveMQTestBase.createSessionFactory(ServerLocator)` for creating your session factory. If you choose *not* to use this method then use `org.apache.activemq.artemis.tests.util.ActiveMQTestBase.addSessionFactory` to add the factory to the test-suite's internal resource ledger.

## Create configurations

There are numerous methods in

`org.apache.activemq.artemis.tests.util.ActiveMQTestBase` to create a configuration. These methods are named like `create*Config(..)`. Each one creates a slightly different configuration but there is a lot of overlap between them.

In any case, `org.apache.activemq.artemis.core.config.Configuration` is a *fluent* interface so it's easy to customize however you need.

## Look at other test-cases

If you need ideas on how to configure something or test something try looking through the test-suite at other test-cases which may be similar. This is one of the best ways to learn how the test-suite works and how you can leverage the testing infrastructure to test your particular case.



## Code coverage report

### Getting JaCoCo exec files

Before you can generate code coverage report by JaCoCo tool, you need to get data about what lines of code were executed during testing. These information are collected by JaCoCo agent and stored in JaCoCo exec files. All you need to do is run the tests with `jacoco` maven profile.

```
mvn test -Ptests,extra-tests,jacoco
```

### Generate JaCoCo reports

```
mvn verify -Pjacoco-generate-report -DskipTests
```

For generating JaCoCo reports only run the maven build with profile `jacoco-generate-report` as it is shown in the example above. After the command was executed, in directory `target/jacoco-report` you can find reports in HTML and XML formats.

### Merge JaCoCo exec files to one

Since ActiveMQ Artemis is divided into several modules, exec files are generated for each module separately. If you need to merge them together to have all data in one place, you can do it by command below.

```
mvn jacoco:merge -N -Pjacoco
```

# Code Formatting

## Eclipse

Eclipse code formatting and (basic) project configuration files can be found at the `etc/` folder. You should manually copy them *after importing all your projects*:

```
for settings_dir in `find . -type d -name .settings`; do
  \cp -v etc/ide-settings/eclipse/org.eclipse.jdt.* $settings_dir
done
```

Do not use the [maven-eclipse-plugin](#) to copy the files as it conflicts with [m2e](#).

## IDEA

If you completed the step described on [idea instructions](#), and selected the code style accordingly you should be ready to go.

## EditorConfig

For editors supporting [EditorConfig](#), a settings file is provided in `etc/ide-settings/editorconfig.ini`. Copy it to your Artemis top level directory and [name it](#) `.editorconfig`

## Validating releases

### Setting up the maven repository

When a release is proposed a maven repository is staged.

This information was extracted from [Guide to Testing Staged Releases](#)

For examples, the 1.1.0 release had the Maven Repository staged as <https://repository.apache.org/content/repositories/orgapachemq-1066>.

The first thing you need to do is to be able to use this release. The easiest way we have found is to change your maven settings at `~/.m2/settings.xml`, setting up the staged repo.

*file ~/.m2/settings.xml:*

```

<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<settings>
  <profiles>
    <profile>
      <id>apache-artemis-test</id>
      <repositories>

        <repository>
          <id>artemis-test</id>
          <name>Apache Artemis Test</name>
          <url>https://repository.apache.org/content/repositories/orgapact

          <layout>default</layout>
          <releases>
            <enabled>true</enabled>
          </releases>
          <snapshots>
            <enabled>true</enabled>
          </snapshots>
        </repository>
      </repositories>

      <pluginRepositories>

        <pluginRepository>
          <id>artemis-test2</id>
          <name>Apache Artemis Test</name>
          <url>https://repository.apache.org/content/repositories/orgapact

          <releases>
            <enabled>true</enabled>
          </releases>
          <snapshots>
            <enabled>true</enabled>
          </snapshots>
        </pluginRepository>
      </pluginRepositories>
    </profile>
  </profiles>

  <activeProfiles>
    <activeProfile>apache-artemis-test</activeProfile>
  </activeProfiles>
</settings>

```

After you configure this, all the maven objects will be available to your builds.

## Using the examples

The Apache ActiveMQ Artemis examples will create servers and use most of the maven components as real application were supposed to do. You can do this by running these examples after the .m2 profile installations for the staged repository.

Of course you can use your own applications after you have staged the maven repository.

## Notes for Maintainers

Core ActiveMQ Artemis members have write access to the Apache ActiveMQ Artemis repositories and will be responsible for acknowledging and pushing commits contributed via pull requests on GitHub.

Core ActiveMQ Artemis members are also able to push their own commits directly to the canonical Apache repository. However, the expectation here is that the developer has made a good effort to test their changes and is reasonably confident that the changes that are being committed will not break the build.

What does it mean to be reasonably confident? If the developer has run the same maven commands that the pull-request builds are running they can be reasonably confident. Currently the [PR build](#) runs this command:

```
mvn -Pfast-tests -Pextra-tests install
```

However, if the changes are significant, touches a wide area of code, or even if the developer just wants a second opinion they are encouraged to engage other members of the community to obtain an additional review prior to pushing. This can easily be done via a pull request on GitHub, a patch file attached to an email or JIRA, commit to a branch in the Apache git repo, etc. Having additional eyes looking at significant changes prior to committing to the main development branches is definitely encouraged if it helps obtain the "reasonable confidence" that the build is not broken and code quality has not decreased.

If the build does break then developer is expected to make their best effort to get the builds fixed in a reasonable amount of time. If it cannot be fixed in a reasonable amount of time the commit can be reverted and re-reviewed.

## Using the dev profile.

Developers are encouraged also to use the Dev profile, which will activate checkstyle during the build:

```
mvn -Pdev install
```

## Commit Messages

Please ensure the commit messages follow the 50/72 format as described [here](#). This format follows the recommendation from the [official Git book](#).

## Configuring git repositories

Aside from the traditional `origin` and `upstream` repositories committers will need an additional reference for the canonical Apache git repository where they will be merging and pushing pull-requests. For the purposes of this document, let's assume these ref/repo associations already exist as described in the [Working with the Code](#) section:

- `origin` : [https://github.com/\(your-user-name\)/activemq-artemis.git](https://github.com/(your-user-name)/activemq-artemis.git)
- `upstream` : <https://github.com/apache/activemq-artemis>
- Add the canonical Apache repository as a remote. Here we call it `apache` .

```
$ git remote add apache https://gitbox.apache.org/repos/asf/activemq-arte
```

- Add the following section to your `/.git/config` statement to fetch all pull requests sent to the GitHub mirror. We are using `upstream` as the remote repo name (as noted above), but the remote repo name may be different if you choose. Just be sure to edit all references to the remote repo name so it's consistent.

```
[remote "upstream"]
  url = git@github.com:apache/activemq-artemis.git
  fetch = +refs/heads/*:refs/remotes/upstream/*
  fetch = +refs/pull/*/head:refs/remotes/upstream/pr/*
```

## Merging and pushing pull requests

Here are the basic commands to retrieve pull requests, merge, and push them to the canonical Apache repository:

1. Download all the remote branches etc... including all the pull requests.

```
$ git fetch --all
Fetching origin
Fetching upstream
remote: Counting objects: 566, done.
remote: Compressing objects: 100% (188/188), done.
remote: Total 566 (delta 64), reused 17 (delta 17), pack-reused 351
Receiving objects: 100% (566/566), 300.67 KiB | 0 bytes/s, done.
Resolving deltas: 100% (78/78), done.
From github.com:apache/activemq-artemis
 * [new ref]          refs/pull/105/head -> upstream/pr/105
```

2. Checkout the pull request you wish to review

```
$ git checkout pr/105 -B 105
```

3. Rebase the branch against master, so the merge would happen at the top of the current master

```
$ git pull --rebase apache master
```

4. Once you've reviewed the change and are ready to merge checkout `master`.

```
$ git checkout master
```

5. Ensure you are up to date on your master also.

```
$ git pull --rebase apache master
```

6. We actually recommend checking out master again, to make sure you wouldn't add any extra commits by accident:

```
$ git fetch apache
$ git checkout apache/master -B master
```

7. Create a new merge commit from the pull-request. IMPORTANT: The commit message here should be something like: "This closes #105" where "105" is the pull request ID. The "#105" shows up as a link in the GitHub UI for navigating to the PR from the commit message. This will ensure the github pull request is closed even if the commit ID changed due to eventual rebases.

```
$ git merge --no-ff 105 -m "This closes #105"
```

8. Push to the canonical Apache repo.

```
$ git push apache master
```

## Using the automated script

If you followed the naming conventions described here you can use the `scripts/rebase-PR.sh` script to automate the merging process. This will execute the exact steps described on this previous section.

- Simply use:

```
$ <checkout-directory>/scripts/merge-pr.sh <PR number> Message on the PR
```

Example:

```
$ pwd
/checkouts/apache-activemq-artemis

$ ./scripts/merge-PR.sh 175 ARTEMIS-229 address on Security Interface
```

The previous example was taken from a real case that generated this [merge commit on #175](#).

- After this you can push to the canonical Apache repo.

```
$ git push apache master
```

## Use a separate branch for your changes

It is recommended that you work away from master for two reasons:

1. When you send a PR, your PR branch could be rebased during the process and your commit ID changed. You might get unexpected conflicts while rebasing your old branch.
2. You could end up pushing things upstream that you didn't intend to. Minimize your risks by working on a branch away from master.

## Notes:

The GitHub mirror repository (i.e. `upstream`) is cloning the canonical Apache repository. Because of this there may be a slight delay between when a commit is pushed to the Apache repo and when that commit is reflected in the GitHub mirror. This may cause some difficulty when trying to push a PR to `apache` that has been merged on the out-of-date GitHub mirror. You can wait for the mirror to update before performing the steps above or you can change your local master branch to track the master branch on the canonical Apache repository rather than the master branch on the GitHub mirror:

```
$ git branch master -u apache/master
```

Where `apache` points to the canonical Apache repository.

If you'd like your local master branch to always track `upstream/master` (i.e. the GitHub mirror) then another way to achieve this is to add another branch that tracks `apache/master` and push from that branch e.g.

```
$ git checkout master
$ git branch apache_master --track apache/master
$ git pull
$ git merge --no-ff pr/105
$ git push
```



## History

The Apache ActiveMQ Artemis project was started in October 2014. The Artemis code base was seeded with a code donation granted by Red Hat, of the HornetQ project. The code donation process consisted of taking a snapshot of the latest HornetQ code base and contributing this snapshot as an [initial git commit](#) into the Artemis git repository.

The HornetQ commit history is preserved and can be accessed here:

<https://github.com/hornetq/hornetq/tree/apache-donation>

Credit should be given to those developers who contributed to the HornetQ project. The top 10 committers are highlighted here:

- Clebert Suconic
- Tim Fox
- Francisco Borges
- Andy Taylor
- Jeff Mesnil
- Ovidiu Feodorov
- Howard Gao
- Justin Bertram
- Trustin Lee
- Adrian Brock

For more information please visit the [HornetQ GitHub project](#).

## Rebasing original donation

It may be useful to look at the donation history combined with the artemis history. It is the case when eventually looking at old changes.

For that there is a script that will rebase master against the donation branch under master/scripts:

- rebase-donation.sh