

Contributing to Infinispan

Table of Contents

1. The Basics	2
1.1. Pre-requisites	2
1.2. Java compatibility	2
1.3. Issue Management	2
1.3.1. Reporting Issues	2
1.3.2. Versioning Guidelines	3
1.4. Version control	3
1.4.1. Setting up your IDE	3
1.5. Builds	8
1.5.1. Continuous Integration	8
1.6. Testing	8
1.7. Communicating with other Infinispan contributors	8
1.8. Style Requirements	8
1.8.1. Spelling	9
1.8.2. Check-in comments	9
1.9. Logging	9
2. Source Control	12
2.1. Prerequisites	12
2.2. Repositories	12
2.3. Roles	12
2.3.1. Contributor	12
2.3.2. Project Admin	16
2.3.3. Release branches	18
2.3.4. Topic branches	18
2.3.5. Comments	19
2.3.6. Commits	19
2.4. Keeping your repo in sync with upstream	19
2.4.1. If you have cloned upstream	19
2.4.2. If you have forked upstream	20
2.5. Tips on enhancing git	21
2.5.1. Completions	21
2.5.2. Terminal colors	21
2.5.3. Aliases	22
2.5.4. Visual History	22
2.5.5. Visual diff and merge tools	22
2.5.6. Choosing an Editor	23
2.5.7. Shell prompt	23
3. Building Infinispan	24

3.1. Requirements	24
3.2. Maven	24
3.2.1. Quick command reference	25
3.2.2. Publishing releases to Maven	27
3.2.3. The Maven Archetypes	28
4. API, Commons and Core	31
4.1. API	31
4.2. Commons	31
4.3. Core	31
5. Running and Writing Tests	32
5.1. Running the tests	32
5.1.1. Specifying which tests to run	32
5.1.2. Skipping the test run	33
5.1.3. Debugging thread leaks	33
5.1.4. Running tests using @Parameters	34
5.1.5. Enabling TRACE in test logs	34
5.1.6. Running tests with JDK 8	34
5.1.7. Enabling code coverage generation	34
5.2. Test groups	35
5.2.1. Which group should I use?	35
5.3. Test permutations	35
5.3.1. Running permutations manually or in an IDE	36
5.4. The Parallel Test Suite	36
5.4.1. Tips for writing and debugging parallel tests	36
6. Helping Others Out	38
7. Adding Configuration	39
7.1. Adding a property	39
7.2. Adding a group	40
7.3. Don't forget to update the XSD and XSD test	40
7.4. Bridging to the old configuration	41
8. Documentation Guidelines	42
8.1. Working with Documentation Source Files	42
8.2. Style Guide	42
8.3. AsciiDoc Format Reference	42
8.4. General Guidelines	43
8.5. Section IDs	44
8.6. Cross References	44
8.7. Diagrams, Screenshots, and Other Media	44
8.7.1. Including Images	45
8.8. Code Samples	45
9. UI Writing Guidelines	46

10. Infinispan Terminology	48
10.1. Infinispan methods and corresponding verbs	49

Feel like contributing to Infinispan? This guide will help you set up your environment, walk you through best practices, and help you debug, improve and add features to Infinispan.

Chapter 1. The Basics

In this chapter we quickly walk through the basics on contributing; future chapters go into more depth.

1.1. Pre-requisites

Java 11	Infinispan is baselined on Java 8 but needs to be built with Java 11 (we build with OpenJDK)
Maven 3.5	The Infinispan build uses Maven, and you should be using at least Maven 3.5
Git	The Infinispan source code is stored in Git.

1.2. Java compatibility

Infinispan can run with Java 8 or greater. However it **must** be compiled at least with Java 11.

1.3. Issue Management

Infinispan uses JIRA for issue management, hosted on issues.redhat.com. You can log in using your jboss.org username and password.

1.3.1. Reporting Issues

If you need to create a new issue then follow these steps.

1. Choose between
 - *Feature Request* if you want to request an enhancement or new feature for Infinispan
 - *Bug* if you have discovered an issue
 - *Task* if you wish to request a documentation, sample or process (e.g. build system) enhancement or issue
2. Then enter a *Summary* , describing briefly the problem - please try to be descriptive!
3. You should **not** set *Priority*.
4. Now, enter the version you are reporting an issue against in the *Affects Version* field, and leave the *Fix Version* field blank.
5. In the *Environment* text area, provide as much detail as possible about your environment (e.g. Java runtime and version, operating system, any network topology which is relevant).
6. In the *Description* field enter a detailed description of your problem or request.
7. If the issue has been discussed on the forums or the mailing list, enter a reference in the *Forum Reference* field
8. Finally, hit *Create*

1.3.2. Versioning Guidelines



Only {brandname} contributors should set the `_Fix Version_` field.

When setting the *Fix Version* field for bugs and issues in JIRA, the following guidelines apply: Version numbers are defined as `${major}.${minor}.${micro}.${modifier}`. For example, `4.1.0.Beta1` would be:

major	4
minor	1
micro	0
modifier	Beta1

If the issue relates to a *Task* or *Feature Request*, please ensure that the `.Final` version is included in the *Fixed In* field. For example, a new feature should contain `4.1.0.Beta1`, `4.1.0.Final` if it is new for 4.1.0 and was first made public in beta1. For example, see [ISPN-299](#).

If the issue relates to a bug which affected a previous Final version, then the *Fixed In* field should also contain the Final version which contains the fix, in addition to any Alpha, Beta or CR release. For example, see [ISPN-546](#). If the issue pertains to a bug in the current release, then the Final version should not be in the *Fixed In* field. For example, a bug found in 4.1.0.Alpha2 (but not in 4.1.0.Alpha1) should be marked as fixed in 4.1.0.Alpha3, but not in 4.1.0.Final. For example, see [ISPN-416](#).

1.4. Version control

Infinispan uses [git](#), hosted on [GitHub](#), for version control. You can find the upstream git repository at <https://github.com/infinispan>. To clone the repository:

```
$ git clone git@github.com:infinispan/infinispan.git
```

or to clone your fork:

```
$ git clone git@github.com:YOUR_GIT_USERNAME/infinispan.git
```


1.4.1. Setting up your IDE

Maven supports generating IDE configuration files for easy setup of a project. This is tested on Eclipse, IntelliJ IDEA and Netbeans.


Eclipse

Before we import the project, we need to clone the project as described above.

1. Install the m2eclipse plugin if you have not already installed it. This is bundled with Eclipse from version "Indigo" 3.7. For older versions follow instructions on <http://eclipse.org/m2e/>.
2. Import the Infinispan maven project. Select File → Import in your eclipse workbench. Select the Existing Maven Project importer.



3. Select the root directory of your Infinispan checkout.



4. Select Infinispan modules that you want to import.
5. Finally, from Infinispan 5.0 onwards, annotation processing is used to allow log messages to be internationalized. This processing can be done directly from Eclipse as part of compilation but it requires some set up:
 - a. Open the properties for infinispan-core and locate Annotation Processing
 - b. Tick Enable project specific settings
 - c. Enter `target/generated-sources/annotations` as the `Generated source` directory



6. Code Formatting. From the menu Window → Preferences → Java → Code Style → Formatter. Import [formatter.xml](#)
7. Code templates. From the menu Window → Preferences → Java → Code Style → Code Templates. Import [codetemplates.xml](#)

IntelliJ IDEA

Importing

When you start [IntelliJ IDEA](#), you will be greeted by a screen as shown below:



If you have already obtained a copy of the Infinispan sources via Github (see '*Source Control*'), then follow: *Import Project* → */directory/to/downloaded/sources* . IntelliJ will automatically make use of Maven to import the project since it will detect a `pom.xml` file in the base directory.

If you have not obtained the sources already, you can use the Git integration in IntelliJ IDEA 12. Click on *Check out from Version Control* → *Github*. After entering your Github credentials, you will then be prompted to enter the git repository URL along with the location that you want to check out the source code to.



Compiler settings

From Infinispan 5.0 onwards, annotation processing is used to allow log messages to be internationalized. This processing can be done directly from IntelliJ as part of compilation but it requires some set up:

1. Go to Preferences → Compiler → Annotation Processor" and click on *Enable annotation processing*
2. Add an annotation processor with "Processor FQN Name" as `org.jboss.logging.LoggingToolsProcessor`
3. In "Processed Modules", add all modules except the root and the parent modules.





There can sometimes be issues with the generated logging classes on rebuild (particularly when you switch Git branches). If these issues do crop up then simply run `mvn clean install -DskipTests` on the command line to clear them out.



If you are running a multi-core environment (e.g. quad-core or above) then you can follow the instructions on making use of parallelized compilation in IntelliJ 12. Information on how to do this can be found [here](#).

Code Style

Download the code style JAR file from [here](#) and import this into IntelliJ IDEA.

1.5. Builds

Infinispan uses [Maven](#) for builds. Make sure you have Maven 3 installed, and properly configured.

1.5.1. Continuous Integration

Infinispan uses [TeamCity](#) for continuous integration. TeamCity polls GitHub for updates and runs whenever updates are available. You can check the status of the latest builds [here](#).

1.6. Testing

Infinispan uses [TestNG](#) for unit and functional tests, and all Infinispan tests are run in parallel. For more information see the chapter on the test suite; this chapter gives advice on writing tests which can safely execute in parallel.

1.7. Communicating with other Infinispan contributors

Infinispan contributors use a mix of technologies to communicate. Visit [this page](#) to learn more.

1.8. Style Requirements

Infinispan uses the [K&R code style](#) for all Java source files, with two exceptions:

1. use 3 spaces instead of a tab character for indentations.
2. braces start on the same line for class, interface and method declarations as well as code blocks.

In addition, sure all [new line characters](#) used must be LF (UNIX style line feeds). Most good IDEs allow you to set this, regardless of operating system used.

All patches or code committed must adhere to this style. Code style settings which can be imported into IntelliJ IDEA and Eclipse are committed in the project sources, in [ide-settings](#).

1.8.1. Spelling

Ensure correct spelling in code, comments, Javadocs, etc. (use *American English* spelling). It is recommended that you use a spellchecker plugin for your IDE.

1.8.2. Check-in comments

Please ensure any commit comments use [this format](#) if related to a task or issue in JIRA. This helps JIRA pick out these checkins and display them on the issue, making it very useful for back/forward porting fixes. If your comment does not follow this format, your commit may not be merged into upstream.

1.9. Logging

Infinispan follows the JBoss logging standards, which can be found [here](#).

From Infinispan 5.0 onwards, Infinispan uses JBoss Logging to abstract over the logging backend. Infinispan supports localization of log message for categories of INFO or above as explained in [the JBoss Logging guidelines](#). As a developer, this means that for each INFO, WARN, ERROR and FATAL message your code emits, you need to modify the Log class in your module and add an explicit method for it with the right annotations.

For example:

```
@LogMessage(level = INFO)
@Message(value = "An informative message: %s - %s", id = 600)
void fiveTransactionsHaveCompleted(String param1, String param2);
```

And then, instead of calling `log.info(...)`, you call the method, for example `log.fiveTransactionsHaveCompleted(param1, param2)`. If what you're trying to log is an error or similar message and you want an exception to be logged as cause, simply use `@Cause` annotation:

```
@LogMessage(level = ERROR)
@Message(value = "An error message: %s", id = 600)
void anErrorMessage(String param1, @Cause IllegalStateException e);
```

The last thing to figure out is which id to give to the message. Each module that logs something in production code that could be internationalized has been given an id range, and so the messages should use an available id in the range for the module where the log call resides. Here are the id range assignments per module:

Module name	Id range
core	1 - 1000
tree	2001 - 3000
bdbje cache store	2001 - 3000

Module name	Id range
cassandra cache store	3001 - 4000
hotrod client	4001 - 5000
server core	5001 - 6000
server hotrod	6001 - 7000
cloud cache store	7001 - 8000
jdbc cache store	8001 - 9000
jdbm cache store	9001 - 10000
remote cache store	10001 - 11000
server memcached	11001 - 12000
server rest	12001 - 13000
server websocket	13001 - 14000
query	14001 - 14800
query-dsl	14801 - 15000
lucene directory	15001 - 16000
<no longer used>	16001 - 17000
cdi integration	17001 - 18000
hbase cache store	18001 - 19000
cli interpreter	19001 - 20000
cli client	20001 - 21000
mongodb cache store	21001 - 22000
rest cache store	22001 - 23000
leveldb cache store	23001 - 24000 [unused]
couchbase cache store	24001 - 25000
redis cache store	25001 - 26000
extended statistics	25001 - 26000
infinispan directory provider	26001 - 27000
tasks	27001 - 27500
scripting	27501 - 28000
remote query server	28001 - 28500
object filter	28501 - 29000
soft-index file store	29001 - 29500
clustered counter	29501 - 30000



When editing the above table, remember to update the README-i18n.txt file in the project sources!



You will need to enable annotation processing in order to be able to compile Infinispan and have the logger implementation generated.

Chapter 2. Source Control

In this chapter we discuss how to interact with Infinispan's source control repository.



As a convention, *upstream* is used as the name of the <http://github.com/infinispan/infinispan> repository. This repository is the canonical repository for Infinispan. We usually name *origin* the fork on [GitHub](#) of each contributor. So the exact meaning of *origin* is relative to the developer: you could think of *origin* as your own fork.

2.1. Prerequisites

This document assumes some working knowledge of git. We recommend Scott Chacon's excellent [Pro Git](#) as a valuable piece of background reading. The book is released under the Creative Commons license and can be downloaded in electronic form for free. At very least, we recommend that you read [Chapter 2](#), [Chapter 3](#) and [Chapter 5](#) of *Pro Git* before proceeding.

2.2. Repositories

Infinispan uses <http://github.com/infinispan/infinispan> as its canonical repository, and this repository contains the stable code on master as well as the maintenance branches for previous minor versions (e.g., 5.0.x, 4.2.x, 4.1.x, etc)

Typically, only *Project Admins* would be able to push to this repository while all else may clone or fork the repository.

2.3. Roles

The project assumes one of 2 *roles* an individual may assume when interacting with the Infinispan codebase. The 2 roles here are:

Roles

- *Contributor*
- *Project Admin* (typically, no more than a small handful of people)



None of the roles assume that you are a Red Hat employee. All it assumes is how much responsibility over the project has been granted to you. Typically, someone may be in more than one role at any given time, and puts on a different "hats" based on the task at hand.

2.3.1. Contributor


A contributor will only ever submit patches via GitHub's *pull request* mechanism.

Contributors should *always* fork the upstream project on GitHub and work off a clone of this fork.



All Infinispan core developers are considered contributors and work off personal forks of the upstream repository. This allows for complex features to be developed in parallel without tripping up over one another. This process is certainly not restricted to just Infinispan core developers; any contributor should also participate in this manner.

Creating a pull request on GitHub



In this workflow, the contributor forks the Infinispan upstream repository on GitHub, clones their fork, and makes changes to this private fork. When changes have been tested and are ready to be contributed back to the project, a *pull request* is issued via GitHub so that one of the *Project Administrators* can pull in the change.

Topic Branches

NOTE: It is desirable to work off a *topic branch*, even when using your own, forked repository. A topic branch is created for every feature or bug fix you do. Typically you would create one topic branch per issue, but if several patches are related it's acceptable to have several commits in the same branch; however different changes should always be identified by different commits.

Before you push your work onto your fork of the repository on GitHub (your *origin*), it is often a good idea to review your commits. Consolidating them (squashing) or breaking them up as necessary and cleaning up commit messages should all be done while still working off your local clone. Also, prior to issuing a pull request, you should make sure you rebase your branch against the upstream branch you expect it to be merged into. Also, only submit pull requests for your topic branch - not for your master!



The section on *Public Small Project* in [Chapter 5, Section 2](#) of Pro Git has more information on this style of workflow.

A worked example

1. Make sure your master is synced up with upstream. See [this section](#) for how to do this
2. Create new branch for your topic and switch to it. For the example issue, ISPN-1234:

```
git checkout -b t_ISPN-12345 master
```

3. Do your work. Test. Repeat
4. Commit your work on your topic branch
5. Push your topic branch to GitHub. For example:

```
git push origin t_ISPN-12345
```

6. Issue a pull request using the [GitHub pull request system](#)
7. Once your pull request has been applied upstream, delete the topic branch both locally and on your fork. For example:

```
git branch -d t_ISPN-12345 && git push origin :t_ISPN-12345
```

8. Sync with upstream again so that your changes now appear in your master branch

If your topic branch has been open for a while and you are afraid changes upstream may clash with your changes, it makes sense to rebase your topic branch before you issue a pull request. To do this:

1. Sync your master branch with upstream

```
git checkout master  
git pull upstream master
```

2. Switch to your topic branch. For example:

```
git checkout t_ISPN-12345
```

3. Rebase your topic branch against master:

```
git rebase master
```

4. During the rebase process you might need to fix conflicts
5. When you're done test your code again.
6. Push your rebased topic branch to your repo on GitHub (you will likely need to force this with the -f option).

```
git push -f origin ISPN-12345
```

7. Continue your work on your topic branch.



If you are sharing your forked Infinispan repo with others, then do not rebase! Use a merge instead.

Multi-step coordination between developers using forked repositories

Sometimes a feature/task is rather complex to implement and requires competence from multiple areas of the projects. In such occasions it is not uncommon for developers to coordinate feature implementation using personal forks of Infinispan prior to finally issuing request to integrate into Infinispan main repository on GitHub.

For example, developer A using his personal Infinispan fork creates a topic branch T and completes as much work as he/she can before requesting for assistance from developer B. Developer A pushes topic T to his personal Infinispan fork where developer B picks it up and brings it down to his local repo. Developer B then in turn completes necessary work, commits his/her changes on branch T, and finally pushes back T to his own personal fork. After issuing request for pull to developer A, developer B waits for notification that developer A integrated his changes. This exchange can be repeated as much as it is necessary and can involve multiple developers.

A worked example

This example assumes that developer A and B have added each others Infinispan forked repositories with the `git add remote` command. For example, developer B would add developer A's personal Infinispan fork repository with the command

```
git remote add devA https://github.com/developerA/infinispan.git
```

1. Developer A starts implementing feature ISPN-244 and works on a local topic branch `t_ISPN244`. Developer A pushes `t_ISPN244` to personal Infinispan fork. For example:

```
git push origin t_ISPN244
```

2. Developer B fetches branch `t_ISPN244` to local repository. For example:

```
git fetch devA t_ispn244:my_t_ispn244
```

3. Developer B works on local branch `my_t_ispn244`
4. Developer B commits changes, pushes `my_t_ispn244` to own fork.

```
git push origin my_t_ispn244
```

5. Developer B sends pull request to developer A to integrate changes from `my_t_ispn244` to `t_ispn244`

2.3.2. Project Admin

Project Admins have a very limited role. Only Project Admins are allowed to push to upstream, and Project Admins *never* write any code directly on the upstream repository. All Project Admins do is pull in and merge changes from contributors (even if the "contributor" happens to be themselves) into upstream, perform code reviews and either commit or reject such changes.



All Contributors who are also Project Admins are encouraged to not merge their own changes, to ensure that all changes are reviewed by someone else.

This approach ensures Infinispan maintains quality on the main code source tree, and allows for important code reviews to take place again ensuring quality. Further, it ensures clean and easily traceable code history and makes sure that more than one person knows about the changes being performed.

Handling pull requests



Project Admins are also responsible for responding to pull requests. When pulling in changes from a forked repository, more than a single commit may be pulled in. Again, this should be done on a newly created working branch, code reviewed, tested and cleaned up as necessary.

If commits need to be altered - e.g., rebasing to squash or split commits, or to alter commit messages - it is often better to contact the Contributor and ask the Contributor to do so and re-issue the pull request, since doing so on the upstream repo could cause update issues for other contributors later on. If commits were altered or three-way merge was performed during a merge instead of fast-forward, it's also a good idea to check the log to make sure that the resulting repository history looks OK:

```
$ git log --pretty=oneline --graph --abbrev-commit # History messed up due to a bad
merge
* 3005020 Merge branch 'ISPN-786' of git://github.com/Sanne/infinispan
|\
| * e757265 ISPN-786 Make dependency to log4j optional <-- Same with cb4e5d6 -
unnecessary
* | cb4e5d6 ISPN-786 Make dependency to log4j optional <-- Cherry-picked commit by
other admin
|/
* ...

$ git reset cb4e5d6 # revert the bad merge
```

It is therefore *strongly recommended* that you use the `handle_pull_request` script that ensures a clean merge. If you *still* wish to do this manually, please consider reading through the script first to get an idea of what needs to happen.



More information on pulling changes from remote, forked repos can be found in [Chapter 5, Section 3](#) of Pro Git, under *Checking Out Remote Branches*.

Possible trouble handling pull requests

1. If you have warnings about "Merge made by recursive" you have to fix it rebasing.
2. If you have warnings about "non-fast-forward" you have to rebase.
3. If you see "non-fast-forward updates were rejected" you **must never** use `--force` on upstream! It means that another patch was merged before you and you have to update your master again, and rebase again.
4. `--force` is allowed only in special maintenance circumstances. If you find you're needing it to handle a pull request, then you're doing it wrong, and the mistake might be a dangerous one! It's like the good rule of never commit when you're drunk (drunk coding, however, is allowed).



Never use `--force` on `git push`

Using `--force` while pushing on a shared repository such as *upstream* you could effectively erase other committed patches. No one should ever use this option unless unanimously approved on the public mailing list: the most dangerous aspect of it is that nobody gets any notification if this happens, and we might think issues are solved but you silently removed the fix and it's history from the repository.

Cutting releases

Releases can only be cut by Project Admins, and must be done off a recently updated (`git fetch` and `git pull origin`) clone of the upstream repo. Infinispan's `bin/release.py` script takes care of the rest.

2.3.3. Release branches

Infinispan has several main release branches. These are master (ongoing work on the current unstable release), and maintenance branches for previous minor releases (e.g., 5.0.x, 4.2.x, 4.1.x). Work should never be committed directly to any of these release branches directly; topic branches should always be used for work, and these topic branches should be merged in using the process outlined above.

2.3.4. Topic branches

Some of the biggest features of git are speed and efficiency of branching, and accuracy of merging. As a result, best practices involve making frequent use of branches. Creating several topic branches a day, even, should not be considered excessive, and working on several topic branches simultaneously again should be commonplace.

[Chapter 3, Section 4](#) of Pro Git has a detailed discussion of topic branches. For Infinispan, it makes sense to create a topic branch and name it after the JIRA it corresponds to. (if it doesn't correspond to a JIRA, a simple but descriptive name should be used).

Topic Branches Affecting More Than One Release Branch

Most topic branches will only affect a single release branch, e.g. features targeted at the current unstable release will only affect the master release branch. So a topic branch should be created based off master. However, occasionally, fixes may apply to both release branches 4.2.x as well as master. In this case, the following workflow should apply:

1. Create topic branch off 4.2.x. For example:

```
git checkout -b <topic>_4.2.x 4.2.x
```

2. Create topic branch off master. For example:

```
git checkout -b <topic>_master master
```

3. Do your work on <topic>_master, test and commit your fixes
4. Switch to <topic>_4.2.x. For example:

```
git checkout <topic>_4.2.x
```

5. Cherry-pick your commit from <topic>_master onto <topic>_4.2.x. For example:

```
git cherry-pick <commit_id>
```

6. Test <topic>_4.2.x for correctness, modify as necessary
7. Issue two separate pull requests for both branches

2.3.5. Comments

It is *extremely important* that comments for each commit are clear and follow certain conventions. This allows for proper parsing of logs by git tools. Read [this article](#) on how to format comments for git and adhere to them. Further to the recommendations in the article, the short summary of the commit message should be in the following format:

ISPN-XXX Subject line of the JIRA in question

This can optionally be followed by a detailed explanation of the commit. Why it was done, how much of it was completed, etc. You may wish to express this as a list, for example:

- Add a unit test
- Add more unit tests
- Fix regressions
- Solve major NP-Complete problems

Make sure however to split separate concerns - especially if they are unrelated - in separate commits.

2.3.6. Commits

Sometimes work on your topic branch may include several commits. For example, committing a test. Then committing another test. Then perhaps committing a fix. And perhaps fixing your own fix in the next commit... Before issuing a pull request for this topic branch, consider cleaning up these commits. Interactive rebasing helps you squash several commits into a single commit, which is often more coherent to deal with for others merging in your work. [Chapter 6, Section 4](#) of Pro Git has details on how to squash commits and generally, clean up a series of commits before sharing this work with others. Note that you can also easily reorder them, just change the order of lines during the interactive rebase process.

Also, it is important to make sure you don't accidentally commit files for which no real changes have happened, but rather, whitespace has been modified. This often happens with some IDEs. `git diff --check` should be run before you issue such a pull request, which will check for such "noise" commits and warn you accordingly. Such files should be reverted and not be committed to the branch.

Adhering to [Infinispan's code style](#) guidelines will help minimise "noise" commits. Project Admins are going to ask contributors to reformat their code if necessary.

2.4. Keeping your repo in sync with upstream

2.4.1. If you have cloned upstream

If you have a clone of the upstream, you may want to update it from time to time. Running:

```
$ git fetch origin
$ git fetch origin --tags
```

will often do the trick. You could then pull the specific branches you would need to update:

```
$ git checkout master
$ git pull origin master
$ git checkout 4.2.x
$ git pull origin 4.2.x
```

Updating topic branches

You should rebase your topic branches at this point so that they are up-to-date and when pulled by upstream, upstream can fast-forward the release branches:

```
$ git checkout <topic>_master
$ git rebase master
```

and/or

```
$ git checkout topic_4.2.x
$ git rebase 4.2.x
```

2.4.2. If you have forked upstream

If you have a fork of upstream, you should probably define upstream as one of your remotes:

```
$ git remote add upstream git://github.com/infinispan/infinispan.git
```

You should now be able to fetch and pull changes from upstream into your local repository, though you should make sure you have no uncommitted changes. (You *do* use topic branches, right?)

```
$ git fetch upstream
$ git fetch upstream --tags
$ git checkout master
$ git pull upstream master
$ git push origin master
$ git checkout 4.2.x
$ git pull upstream 4.2.x
$ git push origin 4.2.x
```



A script can do this for you - have a look at [sync_with_upstream](#).

Updating topic branches

Again, you should rebase your topic branches at this point so that they are up-to-date and when pulled by upstream, upstream can fast-forward the release branches:

```
$ git checkout topic_master  
$ git rebase master
```

and/or

```
$ git checkout topic_4.2.x  
$ git rebase 4.2.x
```

The `sync_with_upstream` script can do this for you if your topic branch naming conventions match the script.

2.5. Tips on enhancing git

2.5.1. Completions

Save [this script](#) as `~/.git-completion.bash` and in `~/.bash_profile`, add the following on one line:

```
source ~/.git-completion.bash
```

After logging out and back in again, typing `git` followed by `TAB` will give you a list of git commands, as would `git c` followed by `TAB`, etc. This even works for options, e.g. `git commit --` followed by `TAB`. The completions are even aware of your refs, so even `git checkout my_br` followed by `TAB` will complete to `git checkout my_branch`!



You get git autocompletion for free if you use `zsh` instead of `bash`.

2.5.2. Terminal colors

Add the following to your `~/.gitconfig`

~/.gitconfig

```
[color]
  ui = yes
[color "branch"]
  current = yellow reverse
  local = yellow
  remote = green
[color "diff"]
  meta = yellow bold
  frag = magenta bold
  old = red bold
  new = green bold
[color "status"]
  added = yellow
  changed = green
  untracked = cyan
```

2.5.3. Aliases

Some git commands are pretty long to type, especially with various switches. Aliases help you to map shortcuts to more complex commands. Again, For example, add the following to `~/.gitconfig`:

~/.gitconfig

```
[alias]
  co = checkout
  undo = reset --hard
  cb = checkout -b
  br = branch
  cp = cherry-pick
  st = status
  l = log --pretty=oneline --decorate --abbrev-commit
  lg = log --decorate --abbrev-commit
  last = log --decorate -1 -p --abbrev-commit
  ci = commit -a
  pom = push origin master
  graph = log --pretty=oneline --graph --abbrev-commit
  dt = difftool
```

2.5.4. Visual History

Git ships with gitk, a GUI that visually represents a log. If you use Mac OS X, [GitX](#) is a good alternative. Try typing gitk or gitx in a git project directory. For Linux users, there are lots of alternatives: *gitk*, *gitg*, *giggle*, ... up to *egit* for Eclipse.

2.5.5. Visual diff and merge tools

There are several options available, including [KDiff3](#), [meld](#) and Perforce's [P4Merge](#) which are all

either open source or available for free. See [this link](#) on setting these up (section under *External Merge and Diff Tools*)

2.5.6. Choosing an Editor

You can customise the editor used by git editing `~/.gitconfig`. The following fires up [MacVIM](#) instead of the default vi editor:

`~/.gitconfig`

```
[core]
  editor = mvim -f
```

Alternatively, you could fire up TextMate or another editors of your choice.

2.5.7. Shell prompt

You can change your bash shell prompt to print the current repository's branch name. Add the following to your `~/.bashrc`

`~/.bashrc`

```
function git_current_branch {
  git branch --no-color 2> /dev/null | sed -e '/^[^*]/d' -e 's/* \(.*)/\1/'
}

if [ "$PS1" ]; then
  PS1='\u@\h:\W]$(git_current_branch)\$ '
fi
```

The resulting shell prompt will look like:

```
trustin@matrix:infinispan-4.2][4.2.x]$
```

If you're a zsh user, you can get even more interesting branch information thanks to [this blog post](#), such as:

- whether your branch is dirty (X)
- whether it's ahead of the remote (↑)
- whether it diverges with the remote (⇅)
- whether it's behind (↓)

For example, the following prompt indicates that the current branch is 't_ispn775_master' and that it is behind remote:

```
[~/Go/code/infinispan.git]% (t_ispn775_master ↓)
```

Chapter 3. Building Infinispan

Infinispan uses [Maven](#) as a build system.

3.1. Requirements

- Java 11 or above
- Maven 3.5 or above



Make sure you follow the steps outlined in [Maven Getting Started - Users](#) to set up your JBoss repository correctly. This step is *crucial* to ensure your Maven setup can locate JBoss artifacts! If you also want to test the EAP integration modules you should also add the appropriate [Enterprise Red Hat Maven Repository](#).

3.2. Maven

The following is an example `settings.xml` to get you started:

settings.xml

```
<settings xmlns="http://maven.apache.org/SETTINGS/1.0.0"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://maven.apache.org/SETTINGS/1.0.0
    http://maven.apache.org/xsd/settings-1.0.0.xsd" >

  <localRepository/>
  <interactiveMode/>
  <usePluginRegistry/>
  <offline/>
  <proxies/>
  <profiles>
    <profile>
      <id>jboss-public-repository</id>
      <repositories>
        <repository>
          <id>jboss-public-repository-group</id>
          <name>JBoss Public Maven Repository Group</name>
          <url> https://repository.jboss.org/nexus/content/groups/public-jboss/ </url>
          <layout>default</layout>
          <releases>
            <enabled>true</enabled>
            <updatePolicy>never</updatePolicy>
          </releases>
          <snapshots>
            <enabled>true</enabled>
            <updatePolicy>never</updatePolicy>
          </snapshots>
        </repository>
      </repositories>
    </profile>
  </profiles>
</settings>
```

```

<pluginRepositories>
  <pluginRepository>
    <id>jboss-public-repository-group</id>
    <name>JBoss Public Maven Repository Group</name>
    <url> https://repository.jboss.org/nexus/content/groups/public-jboss/ </url>
    <layout>default</layout>
    <releases>
      <enabled>true</enabled>
      <updatePolicy>never</updatePolicy>
    </releases>
    <snapshots>
      <enabled>true</enabled>
      <updatePolicy>never</updatePolicy>
    </snapshots>
  </pluginRepository>
</pluginRepositories>
</profile>

<!-- Include early access of application server and other products -->
<profile>
  <id>redhat-earlyaccess-repository</id>
  <repositories>
    <repository>
      <id>redhat-earlyaccess-repository-group</id>
      <name>Red Hat early access repository</name>
      <url> http://maven.repository.redhat.com/earlyaccess/all/ </url>
      <layout>default</layout>
      <releases>
        <enabled>true</enabled>
        <updatePolicy>never</updatePolicy>
      </releases>
      <snapshots>
        <enabled>true</enabled>
        <updatePolicy>never</updatePolicy>
      </snapshots>
    </repository>
  </repositories>
</profile>
</profiles>
<activeProfiles>
  <activeProfile>jboss-public-repository</activeProfile>
  <activeProfile>redhat-earlyaccess-repository</activeProfile>
</activeProfiles>
</settings>

```

3.2.1. Quick command reference



Maven places it's output in **target/**

Command	Meaning
<code>mvn clean</code>	Cleans out any old builds and binaries
<code>mvn compile</code>	Compiles java source code
<code>mvn test</code>	Runs the TestNG unit test suite on the compiled code. Will also compile the tests. See the testing section below for more information on running different test groups. The default test group run is the "unit" group.
<code>mvn package</code>	Packages the module as a JAR file, the resulting JAR file will be in target/
<code>mvn package -DskipTests</code>	Creates a JAR file without running tests
<code>mvn package -DskipTests -P minimal-distribution</code>	Creates a reduced version of the distribution with all modules,scripts...etc but no javadoc or source code. This is very handy to quickly build the distribution in order to run some tests.
<code>mvn install -DskipTests</code>	Installs the artifacts in your local repo for use by other projects/modules, including inter-module dependencies within the same project.
<code>mvn install -P distribution</code>	In addition to install, will also use Maven's assembly plugin to build ZIP files for distribution (in target/distribution). Contents of various distribution are controlled by the files in src/main/resources/assembly .

Command	Meaning
<code>mvn deploy</code>	Builds and deploy the project to the JBoss snapshots repository.
<code>mvn install -P-extras</code>	Avoids the extras profile disables the enforce plugin and generation of source jars, hence making builds run faster. Clearly, this option should not be used when making a release or publishing a snapshot.



For non-snapshot releases (e.g., alphas, betas, release candidates and final releases) you should use the `bin/release.py` script.

3.2.2. Publishing releases to Maven

To be able to publish releases to Maven, you need to have the following in your `settings.xml` file:

settings.xml

```
<settings>
...
<servers>
...
<server>
  <id>jboss-snapshots-repository</id>
  <username>your JBoss.org username</username>
  <password>your JBoss.org password</password>

</server>
<server>
  <id>jboss-releases-repository</id>
  <username>your JBoss.org username</username>
  <password>your JBoss.org password</password>

</server>
...
</servers>
...
</settings>
```

Publishing snapshots

Simply running

```
$ mvn clean deploy -DskipTests
```

in the Infinispan root directory will deploy a snapshot.

Publishing releases

Use the `bin/release.py` script.

3.2.3. The Maven Archetypes

Infinispan currently has 2 separate Maven [archetypes](#) you can use to create a skeleton project and get started using Infinispan. This is an easy way to get started using Infinispan as the archetype generates sample code, a sample Maven pom.xml with necessary dependencies, etc.

You don't need to have any experience with or knowledge of Maven's Archetypes to use this! Just follow the simple steps below.

Starting a new project

Use the `newproject-archetype` project. The simple command below will get you started, and

```
$ mvn archetype:generate \  
  -DarchetypeGroupId=org.infinispan.archetypes \  
  -DarchetypeArtifactId=newproject-archetype \  
  -DarchetypeVersion=1.0.5 \  
  -DarchetypeRepository=http://repository.jboss.org/nexus/content/groups/public
```

You will be prompted for a few things, including the artifactId , groupId and version of your new project. And that's it - you're ready to go!

Exploring your new project

The skeleton project ships with a sample application class for interacting with Infinispan. You can open this new project in your IDE - most good IDEs such as IntelliJ and Eclipse allow you to import Maven projects, see [this guide](#) and [this guide](#). Once you open your project in your IDE, you should examine the generated classes and read through the comments.

On the command line...

Try running

```
$ mvn install -Prun
```

in your newly generated project. This runs the `main()` method in the generated application class.

Writing a test case for Infinispan

This archetype is useful if you wish to contribute a test to the Infinispan project and helps you get set up to use Infinispan's testing harness and related tools. Use

```
$ mvn archetype:generate \  
  -DarchetypeGroupId=org.infinispan.archetypes \  
  -DarchetypeArtifactId=testcase-archetype \  
  -DarchetypeVersion=1.0.5 \  
  -DarchetypeRepository=http://repository.jboss.org/nexus/content/groups/public
```

As above, this will prompt you for project details and again as above, you should open this project in your IDE. Once you have done so, you will see some sample tests written for Infinispan making use of Infinispan's test harness and testing tools along with extensive comments and links for further reading.

On the command line...

Try running

```
$ mvn test
```

in your newly generated project to run your tests. The generated project has a few different profiles you can use as well, using Maven's -P flag. For example:

```
$ mvn test -Pudp
```

Available profiles

The profiles available in the generated sample project are:

- udp: use UDP for network communications rather than TCP
- tcp: use TCP for network communications rather than UDP
- jbosstm: Use the embedded [JBoss Transaction Manager](#) rather than Infinispan's embedded transaction manager

Contributing tests back to Infinispan

If you have written a functional, unit or stress test for Infinispan and want to contribute this back to Infinispan, your best bet is to [fork the Infinispan sources on GitHub](#). The test you would have prototyped and tested in an isolated project created using this archetype can be simply dropped in to Infinispan's test suite. Make your changes, add your test, prove that it fails even on Infinispan's upstream source tree and issue a [pull request](#).

Checking coding style

If you have written any new code, it is highly recommended to validate formatting before submitting a Pull Request. This might be done by invoking:

```
$ mvn checkstyle:check
```

Versions

The archetypes generate poms with dependencies to specific versions of Infinispan. You should edit these generated poms by hand to point to other versions of Infinispan that you are interested in.

Source Code

The source code used to generate these archetypes are [on GitHub](#). If you wish to enhance and contribute back to the project, fork away!

Chapter 4. API, Commons and Core

In order to provide proper separation between public APIs, common utilities and the actual implementation of Infinispan, these are split into 3 Maven modules: `infinispan-api`, `infinispan-commons` and `infinispan-core`. This separation also makes sure that modules, such as the remote clients, don't have to depend on `infinispan-core` and its transitive dependencies. The following paragraphs describe the role of each of these modules and give indication as to what goes where.

4.1. API

The `infinispan-api` module should only contain the public interfaces which can be used in any context (local, remote, etc). Any additions and/or modifications to this module *must* be discussed and approved by the core team beforehand. Ideally it should not contain any concrete classes: rare exceptions may be made for small, self-contained classes which need to be referenced from the API interfaces and for which the introduction of an interface would be deemed cumbersome.

4.2. Commons

The `infinispan-commons` module contains utility classes which can be reused across other modules. Classes in `infinispan-commons` should be self-contained and not pull in any dependencies (apart from the existing `jboss-logging` and `infinispan-api`). They should also make no reference to configuration aspects specific to a particular environment.

4.3. Core

The `infinispan-core` module contains the actual implementation used for local/embedded mode. When adding new functionality to the APIs, it is generally safe to start by putting them in `infinispan-core` and promoting them to `infinispan-api` only when it is deemed mature enough to do so and it makes sense across the various use-cases.

Chapter 5. Running and Writing Tests

Tests are written using the [TestNG](#) framework.

5.1. Running the tests

The default run executes all tests in the functional, unit, xsite, and arquillian groups. To just run the tests with txt and xml output the command is:

```
$ mvn test
```

Alternatively, you can execute the tests *and* generate a report with:

```
$ mvn surefire-report:report
```

If you are running the tests on a Unix-like operating system, the default limits per user are typically low. The Infinispan test suite creates a lot of processes/threads, thus you will have to increase your user's limits and reboot the system to pick up the new values. Open up [/etc/security/limits.conf](#) and add the following lines replacing the user name with your username.

[/etc/security/limits.conf](#)

rhusar	soft	nofile	16384
rhusar	hard	nofile	16384
rhusar	soft	nproc	16384
rhusar	hard	nproc	16384

We recommend running the tests on a machine with at least 4GB of RAM. Tests run in a forked JVM, so [MAVEN_OPTS](#) does not affect them. However, you may need to reduce the Maven JVM's heap to run in 4GB of RAM by updating the [MAVEN_OPTS](#) environment variable, e.g.

```
$ export MAVEN_OPTS="-Xmx800m"
```

5.1.1. Specifying which tests to run

A single test can be executed using the test property. The value is the short name (not the fully qualified package name) of the test. For example:

```
$ mvn -Dtest=FqnTest test
```

Alternatively, if there is more than one test with a given classname in your test suite, you could provide the path to the test.

```
$ mvn -Dtest=org/infinispan/api/MixedModeTest test
```

Patterns are also supported:

```
$ mvn -Dtest=org/infinispan/api/* test
```

Also, you can always pass your own Log4j configuration file via `-Dlog4.configuration` with your own logging settings.

5.1.2. Skipping the test run

It is sometimes desirable to install the Infinispan package in your local repository without performing a full test run. To do this, simply use the `skipTests` property:

```
$ mvn -DskipTests install
```

Note that you should *never* use `-Dmaven.test.skip=true` since modules' test classes depend on other module test classes, and this will cause compilation errors.

5.1.3. Debugging thread leaks

The Infinispan test suite uses TestNG/JUnit test listeners to find and report thread leaks. In order to debug thread leaks in the IDE, please add this listener to your default TestNG run configuration:

```
org.infinispan.commons.test.TestNGTestListener
```

If your IDE supports listeners in the JUnit run configuration (IntelliJ IDEA does not), add this listener:

```
org.infinispan.commons.test.JUnitTestListener
```

If your IDE does not support listeners in the JUnit run configuration, you can always add the listener to your class:

```
@Listeners(org.infinispan.commons.test.JUnitTestListener.class)
```

In order to make the leak report look like a test failure, modules using TestNG duplicate the check in a test, `TestNGSuiteChecksTest`. Including the test is not required to see the leak, but without it the failure looks more like a crash.

5.1.4. Running tests using `@Parameters`

If you want to execute tests relying on TestNG's `@Parameters` from an IDE (such as Eclipse or IntelliJ IDEA), please read [this blog entry](#).

5.1.5. Enabling TRACE in test logs

When you run tests, you can get TRACE logging via using the `traceTests` profile

```
$ mvn test -PtraceTests
```

Executing this will generate a GZIP file called `trace-infinispan.log.gz`. This file is not fully closed, so to extract the log file, execute:

```
$ gunzip -c trace-infinispan.log.gz > trace-infinispan.log
```

5.1.6. Running tests with JDK 8

While building requires at least JDK 11, the testsuite can be run with JDK 8. In order to do so, you should set the `JAVA8_HOME` environment variable to point to your JDK 8 installation, e.g.:

```
$ export JAVA8_HOME=/usr/lib/jvm/java-1.8.0-openjdk-amd64
```

And enable the `java8-test` maven profile:

```
$ mvn -Pjava8-test test
```

5.1.7. Enabling code coverage generation

When you run tests, you can enable code coverage recorder for calculating and analysing the Infinispan code coverage. You can do this using `coverage` and `jacocoReport` profiles. As a code coverage evaluation tool, the JaCoCo is used.

```
$ mvn test -Pcoverage -Dmaven.test.failure.ignore=true
```

Please note, that `-Dmaven.test.failure.ignore=true` is used for generating JaCoCo code coverage report, even if there are test failures.

Executing this will generate `jacoco.exec` file in each module's `target/` directory. These are the JaCoCo execution data files, which contain full data about the specific module's coverage.

As soon as the coverage execution command is finished, you will need to generate the JaCoCo report, which will merge the generated `jacoco.exec` files as well as will create the code coverage report.

For having the report in place, run the following command from your Infinispan home directory:

```
$ mvn install -pl parent -PjacocoReport
```

The `jacoco-html/` directory will be generated in Infinispan Home directory, which will contain the code coverage report.

5.2. Test groups

Each test should belong to one or more group. The group acts as a filter, and is used to select which tests are ran as part of the maven test lifecycle.

5.2.1. Which group should I use?

The following test groups are used by Infinispan.



If your test does not fit into one of these groups, a new group should be added.

Test Group	Description
<i>unit</i>	Unit tests using stubs to isolate and test each major class in Infinispan. This is the default group run if no test group is specified
<i>functional</i>	Tests which test the general functionality of Infinispan
<i>jgroups</i>	Tests which need to send data on a JGroups Channel
<i>transaction</i>	Tests which use a transaction manager
<i>profiling</i>	Tests used for manual profiling, not meant for automated test runs
<i>manual</i>	Other tests that are run manually

Every test (except those not intended to be run by continuous integration) should at least be in the **functional** or **unit** groups, since these are the default test groups executed by Maven, and are run when preparing a release.

5.3. Test permutations

We use the term permutation to describe a test suite execution against a particular configuration. This allows us to test a variety of environments and configurations without rewriting the same basic test over and over again. For example, if we pass JVM parameter `-Dinfinispan.cluster.stack=udp` test suite is executed using UDP config.

```
$ mvn -Dinfinispan.cluster.stack=udp test
```

Each permutation uses its own report directory, and its own html output file name. This allows you to execute multiple permutations without wiping the results from the previous run. Note that due to the way Maven operates, only one permutation can be executed per `mvn` invocation. So automating multiple runs requires shell scripting, or some other execution framework to make multiple calls to Maven.

5.3.1. Running permutations manually or in an IDE

Sometimes you want to run a test using settings other than the defaults (such as UDP for `jgroups` group tests or the `EmbeddedTransactionManager` for `transaction` group tests). This can be achieved by referring to the Maven POM file to figure out which system properties are passed in to the test when doing something different. For example to run a `jgroups` group test in your IDE using TCP instead of the default UDP, set `-Dinfinispan.cluster.stack=tcp`. Or, to use JBoss JTA (Arjuna TM) instead of the `EmbeddedTransactionManager` in a `transaction` group test, set `-Dinfinispan.test.jta.tm=jbosstm`. Please refer to the POM file for more properties and permutations.

5.4. The Parallel Test Suite

Infinispan runs its unit test suite in parallel; Infinispan tests are often IO rather than processor bound, so executing them in parallel offers significant speed up in executing the entire test suite.

5.4.1. Tips for writing and debugging parallel tests

There are a number of constraints and best practices that need to be followed in order to ensure correctness and keep the execution time to a minimum. If you follow these guidelines you will find your tests are more reliable:

- *Each test class is run in a single thread* There is no need to synchronize unit test's fixture, as test methods will be run in sequence. However, multiple test classes are executed in parallel.
- *Each test class must have an unique test name* As a convention, the name of the test should be the fully qualified class name of the test class with the `org.infinispan` prefix removed. For example, given a test class `org.infinispan.mypackage.MyTest` the name of the test should be `mypackage.MyTest`. This convention guarantees a unique name.

MyTest.java

```
package org.infinispan.mypackage;
@Test (testName = "mypackage.MyTest")
public class MyTest { ... }
```

- Use `TestCacheManagerFactory.createXyzCacheManager` and **never** create managers using `new DefaultCacheManager()`. This ensures that there are no conflicts on resources e.g. a cluster created by one test won't interfere with a cluster created by another test.

- Where possible, extend `SingleCacheManagerTestorMultipleCacheManagersTest`. Tests inheriting from these template method classes will only create a cache/cluster once for all the test methods, rather than before each method. This helps keep the execution time down.
- **Never** rely on `Thread.sleep()`. When running in heavily threaded environments this will most often not work. For example, if using `ASYNC_REPL`, don't use a `sleep(someValue)` and expect the data will be replicated to another cache instance after this delay has elapsed. Instead, use a `ReplicationListener` (look up javadocs for more information on this class). Generally speaking, if you expect something will happen and you don't have a guarantee when, a good approach is to try that expectation in a loop, several times, with an generous (5-10secs) timeout. For example:

```
while (System.currentTimeMillis() - startTime < timeout) {
    if (conditionMet()) break;
    Thread.sleep(50);
}
```

- `Thread.sleep(10)` may not work in certain OS/JRE combos (e.g. Windows XP/Sun JRE 1.5). Use values greater than 10 for these statements, e.g. 50. Otherwise, a `System.currentTimeMillis()` might return same value when called before and after such a sleep statement.
- For each cache that is created with `TestCacheManagerFactory.createXyzCacheManager()` the test harness will allocate a unique JMX domain name which can be obtained through `CacheManager.getJmxDomain()`. This ensures that no JMX collisions will take place between any tests executed in parallel. If you want to enforce a JMX domain name, this can be done by using one of the `TestCacheManagerFactory.createCacheManagerEnforceJmxDomain` methods. These methods must be used with care, and you are responsible for ensuring no domain name collisions happen when the parallel suite is executed.
- Use obscure words. Insert uncommon or obscure words into the cache that has been generated with a random word generator. In a multi-threaded environment like Infinispan's testsuite, grepping for these words can greatly help the debugging process. You may find [this random word generator](#) useful.
- Use the test method name as the key. Grab the test method and use it as part of the cached key. You can dynamically grab the test method using code like this:

```
Thread.currentThread().getStackTrace()(1).getMethodName()
```



Even though we've tried to reduce them to a minimum, intermittent failures might still appear from time to time. If you see such failures *in existing code* please report them in the issue tracker.

Chapter 6. Helping Others Out

Infinispan is reliant on the whole community helping each other out. Less experienced contributors are often able to help out answering the "newbie" questions, leaving more experienced contributors to handle the more complex questions.

Users are encouraged to follow the [How to ask a forum question](#) guide.

Forum discussions should be posed as questions (we encourage people to do this). Questions can be marked as answered, indicating to the community that they no longer require answering, allowing easy tracking of open questions. Open questions can [be easily viewed using this filter](#). Community members are encouraged to regularly view the open questions and answer any questions they can.

In order to track what questions are still open, you are encouraged to mark questions as "assumed answered" if you provide information that you think resolves the query and you don't hear back to the contrary after a week or so.

Approximately every month, a member of the Infinispan team will go through any open questions for the past month and clear up any unanswered questions, either by chasing for an answer from core team, or by checking with the user if they require more info.

Chapter 7. Adding Configuration



We still use the old configuration system internally within Infinispan. This makes things a little complicated. This will be switched out soon! For now, you need to also add your property to the old config system as well as the new.

Note, these guides assume you are adding an element to the cache configuration, but apply equally to the global configuration.

Before you start adding a configuration property, identify whether you want to add a property to an existing configuration group/element, or whether you need to create a child object. We call the configuration group XXX in the steps below.

7.1. Adding a property

1. Add the property to the relevant XXXConfiguration class. Add a private final field, add a parameter to the constructor, and assign the value to the field in the constructor body. Add an accessor for the property. If the property should be mutable at runtime, then add a mutator as well. Most configuration properties are not mutable at runtime - if the configuration is runtime mutable, then Infinispan needs to take notice of this update whilst the cache is running (you can't cache the value of the configuration in your implementation class). Mutators and accessors don't use the classic JavaBean pattern of prepending accessors with "get" and mutators with "set". Instead, the name of the property is used for an accessor. A mutator is an overloaded version of the accessor which takes a parameter, the new value.
2. Add the property to the matching XXXConfigurationBuilder. You'll need to add a mutable field to the class, and initialise it to its default value in the field declaration. Add a mutator (following the above pattern).
3. The `create()` method is called by the parent object in order to instantiate the XXXConfiguration object from the builder. Therefore, make sure to pass the value of the field in the builder to the XXXConfiguration object's constructor here. Additionally, if you require a complex default (for example, the value of a configuration property is defaulted conditionally based on the value of some other configuration property), then this is the place to do this.
4. The `validate()` method is called by the parent object to validate the values the user has passed in. This method may also be called directly by user code, should they wish to manually validate a configuration object. You should place any validation logic here related to your configuration property. If you need to "cross-validate" properties (validate the value of your property conditionally upon the value of another property), and the other property is on another builder object, increase the visibility of that other property field to "default", and reference it from this builder, by calling the `getBuilder()` method, which will give you a handle on the root configuration builder.
5. The final step is to add parsing logic to the `Parser` class. First, add the attribute to name to the `Attribute` enum (this class simply provides a mapping between the non-type-safe name of the attribute in XML and a type-safe reference to use in the parser). Locate the relevant `parseXXX()` method on the class, and add a case to the switch statement for the attribute. Call the builder mutator you created above, performing any XML related validation (you are unlikely to need

this), and type conversion (using the static methods on the primitive wrapper classes, String class, or relevant enum class).

7.2. Adding a group

In some situations you may additionally want to add a configuration grouping object, represented in XML as an element. You might want to do this if you are adding a new area of functionality to Infinispan. Identify the location of the new configuration grouping object. It might be added to the root Configuration object, or it might be added to one of its children, children's children. We'll call the parent `YYY` in the steps below. Create the `XXXConfiguration` object. Add any properties required following the guide for adding properties. The constructor's visibility should be "default".

1. Create the `XXXConfigurationBuilder` object. It should subclass the relevant configuration child builder – use the `YYYConfigurationChildBuilder` as the superclass. This will ensure that all builder methods that allow the user to "escape" are provided correctly (i.e. provide access to other grouping elements), and also require you to provide a `create()` and `validate()` method. The constructor needs to take the `YYYConfigurationBuilder` as an argument, and pass this to the superclass (this simply allows access to the root of the builder tree using the `getBuilder()` method).
2. Follow the property adding guide to add any properties you need to the builder. The `create()` method needs to return a new instance of the `XXXConfiguration` object. Implement any validation needed in the `validate()` method.
3. In the `YYYConfiguration` object, add your new configuration class as a private final field, add an accessor, and add initialiser assignment in the constructor

• In the `YYYConfigurationBuilder`, add your new configuration builder as a private final field, and initialise it in the constructor with a new instance. Finally, add an accessor for it following the standard pattern discussed in the guide.

1. In the `YYYConfigurationBuilder` ensure that your `validate` method is called in its `validate` method, and that result of the `XXXConfiguration` instances' `create` method is passed to the constructor of `YYYConfiguration`.
2. Finally, add this to the parser. First, add the element to the `Element` class, which provides a type safe representation of the element name in XML. In the `Parser` class, add a new `parseXXX` method, copying one of the others that most matches your requirements (parse methods either parse elements only - look for `ParseUtils.requireNoAttributes()`, attributes only – look for `ParseUtils.requireNoContent()` or a combination of both – look for an iterator over both elements and attributes). Add any attributes as discussed in the adding a property guide. Finally, wire this in by locating the `parseYYY()` method, and adding an element to the switch statement, that calls your new `parseXXX()` method.

7.3. Don't forget to update the XSD and XSD test

1. Add your new elements or attributes to the XSD in `src/main/resources`.
2. Update `src/test/resources/configs/all.xml` to include your new elements or attributes.

7.4. Bridging to the old configuration

Until we entirely swap out the old configuration you will need to add your property to the old configuration (no need to worry about JAXB mappings though!), and then add some code to the `LegacyConfigurationAdaptor` to adapt both ways. It's fairly straightforward, just locate the relevant point in the `adapt()` method (near the configuration group you are using) and map from the legacy configuration to the new configuration, or vs versa. You will need to map both ways, in both adapt methods.

Chapter 8. Documentation Guidelines

Welcome to Infinispan documentation. This is a place where anyone can contribute and share their knowledge.

8.1. Working with Documentation Source Files

All you need to do is fork Infinispan to edit the documentation source files. You can then submit your changes in a *pull request* as you would any other contribution to the project.

See the [Documentation README](#) for information on what you need to edit and build documentation.

8.2. Style Guide

Infinispan is the community project on which Red Hat Data Grid is based. For this reason, we follow [Guidelines for Red Hat Documentation](#).

The [AsciiDoctor Writer's Guide](#) also has some excellent resources on writing effective documentation.

8.3. AsciiDoc Format Reference

Tips for which AsciiDoc markup to use when formatting text.

Monospace (``)

Item	Example
File names and paths	The <code>`/home/user/.bashrc`</code> file ...
Literal values	If the value is <code>`true`</code> ...
Configuration attributes	Set the <code>`enabled`</code> attribute ...
Java class names	The <code>`java.sql.Connection`</code> class ...
Java annotations	The <code>`@Clustered`</code> annotation ...

Italics (_)

Item	Example
Guide titles	See the <code>_Installation Guide_</code> ...
Emphasis (<i>first occurrence of the term only</i>)	<code>_High availability_</code> refers to the ...

Bold (*)

Item	Example
GUI items (<i>menus, buttons</i>)	Click the <code>*Add*</code> button and ...

Underline (value)

Item	Example
Default item in a multi-select	`yes\[.underline]#no#\maybe`

8.4. General Guidelines

- Always use `{brandname}` to refer to the product. Never abbreviate the product name to something like "ISPN".
- Prefer active voice to passive voice.
- Use the second or third person, but never the first. Prefer second person singular to address the user directly.
- Always use the present tense. There are instances where future tense makes sense but do not mix tenses. Never use the past tense.
- Use American English spelling conventions.
- Avoid contractions. Use "do not" instead of "don't".
- Terminate bulleted and numbered lists with periods (.) except for simple name lists.
- Avoid the word "simply" unless it clarifies.
- For substitution of `{attr}` in code blocks, use `[subs=+attributes]`.
- For styling of `*bold*` (**bold**) in code blocks, use `[subs=+quotes]`.
- Never use colloquialisms or slang in documentation.

Item	Use	Do Not Use
Root commands	\$ sudo	#
Emphasis	_yay_	*yay*
Decimal integers < 10	four	4
Decimal integers >= 10	12	twelve
Hexadecimal integers (<i>always numeric, lowercase x</i>)	0x123	0X13
Number ranges (<i>always use numerals</i>)	1-20	1-twenty
Avoid Latin abbreviations	that is	i.e.
	and so on	etc.
	for example	e.g.

Terminology

See the [Infinispan Terminology](#) section for guidance on using words and phrases specific to Infinispan capabilities and features.

One Sentence Per Line

Rather than using a soft limit of characters per line, put one sentence on each line.

Notes and Admonitions

Encapsulate notes, tips, and warnings in brackets. Surround the text in four equals characters, as follows:

```
[NOTE]
====
This is a note that you should read.
====
```

8.5. Section IDs

When creating anchors for section IDs, use the following format:

```
[[anchor_name]]
```

You should always use underscores (`_`) to delimit anchors. Do not use other characters such as dashes (`-`) or periods (`.`) to avoid issues at build time.

To create anchors for reusable content, use this format:

```
[id="same-as-section-heading-{context}"]
```

Learn more about content reuse in the [Modular Documentation Reference Guide](#).

8.6. Cross References

To reference a section within the same book, use `link:#`:

```
link:#anchor_name[Link Text]
```

8.7. Diagrams, Screenshots, and Other Media

- Images should be saved as **PNG** or **JPG**, with a width of at least **660 px**, at **110 dpi**. Try to keep file size less than **300 KB**.
- Screenshots supplement the text, not replace it. **Do not use images as the sole means to convey information or context.**
- **Do not include any test or pre-release labels.**
- **Do not include any personally identifying information.**
- Capture just the part of the screen or window that users must focus on; **do not include window**

headers in the final screenshots unless completely necessary.

- Crop screenshots to **condense important information** and limit empty GUI space and other inconsequential parts.
- All information in an image must be available in an alternative text format for accessibility (Section 508).
- Save all images under `documentation/src/main/asciidoc/${your_document}/images`

8.7.1. Including Images

Insert images using the `image::` or `image:` directive.

- Example 1: Image title in title case (which automatically appends a Figure #).

```
.Image Title  
image::icon.png[Alt text, 50, 50]
```

- Example 2: Inline image. Note, there is only one colon (:) used here.

```
This is an inline image. image:icon.png[Alt text] Cool!
```

8.8. Code Samples

Include code samples in blocks such as the following:

```
[source,java,options="nowrap"]  
.MyClass.java  
----  
//some Java code  
----
```



Include code samples that demonstrate an idea. To share reusable blocks of code or configuration files, store them in GitHub as a [gist](#) and linking to them.

Chapter 9. UI Writing Guidelines

This section describes guidelines for working with text in Infinispan user interface (UI) elements and applies, but is not limited, to the following components:

- Infinispan Console
- Infinispan CLI
- log messages
- error messages

Content in the Infinispan UI should adhere to the [PatternFly UX writing style guide](#).

While the PatternFly design guidelines are the definitive resource for UI writing guidelines, the following sections outline some rules for creating microcopy in Infinispan UI.

Headings

Headings on UI elements should use sentence casing. See the [PatternFly capitalization rules](#)

Confirmation dialogs

Dialogs that prompt for user action should follow this format:

- **Headline:** A headline is usually phrased as a question. Include keywords (like “permanent”) in the headline.
- **Body text:** Body text gives information about the action’s consequence.
- **Buttons:** Buttons allow a user to answer the headline question.



Confirmation dialogs disrupt user task flow. For this reason, you should not provide confirmation dialogs for actions that are easily reversed or insignificant.

Success or failure messages

Use simple past tense form for success or failure messages.

Correct: "Cache created." **Incorrect:** "Cache has been created."

Success or failure messages do not need to be complete sentences. If your message is a fragment, do not include a period/full stop character (.).

Tooltip and hover help

Icons and buttons with action should either have labels or tooltips.

For accessibility reasons, you should add tooltips only to semantic icons that do not have labels.

See the [PatternFly accessibility guide](#).

Abbreviations

Use common and familiar abbreviations only, such as "URL". Write out all other abbreviations, especially the following:

- Number: Avoid "No."
- Average: Avoid "avg."
- For example: Avoid "e.g."
- In other words: Avoid "i.e."

For more information on abbreviations and acronyms, see the Numerics page in the [PatternFly UX writing style guide](#).

Chapter 10. Infinispan Terminology

Infinispan terminology refers to words and phrases specific to Infinispan capabilities and features. It is important to use consistent terminology across documentation and the code base to avoid unclear or contradictory language.

Add

Add an item to an existing list, group, view, or other container element. For example, "add entries to caches". Opposite to "Remove".

Cache

A data structure for storing key/value pairs in memory.

Cache configuration

A resource that instantiates and controls a cache.

Cache template

Generally means a resource from which you create cache configurations. Infinispan provides some pre-defined templates (`org.infinispan.*`). Users can define new templates in `infinispan.xml` with the `*-cache-configuration*` tags. In some APIs, you can use the configuration of an existing cache as a template to configure a new cache.

Cache definition

Refers to a named cache instance. To avoid ambiguity, do not use in UI or documentation to mean "cache configuration".

Cache Manager

Interface for creating and managing caches. Always spell as two words with capital letters when referring to the abstract notion of a Cache Manager. When referring to specific interfaces, use `CacheManager`, `EmbeddedCacheManager`, or `RemoteCacheManager`.

Clear

Delete all the elements of a container. For example, "clear entries in a cache".

Create

Build something new. For example, "create a cache". Opposite to "Delete".

Delete

Permanently remove a single element or a group of similar elements from a container. For example, "delete an entry from the cache" or "delete all expired entries from the cache". Opposite to "Create".

Destroy

Do not use to mean "delete".

Index

Create an index to improve query performance or add an entry to an existing index when it is added to a cache.

Indexes

Plural of index. Do not use "indices".

Off-heap

Native memory space outside the managed JVM heap. Always hyphenate when used as an adjective.

Please

Avoid using the word "please" because it can imply that actions are optional. Technical documentation should be authoritative and clear. Saying "please" can make language sound more "human" but politeness is not necessary in technical content.

Purge

Can have negative connotations and is not a preferred term. Use only as a synonym for "clear" or "delete" in subsequent descriptions, such as in the body text of a dialog or text that follows a heading. Do not use as a primary verb or in headings.

Query

Look up information within the data set. The words "query" and "search" are synonymous but frequently interchanging the two can lead to inconsistency. Use in verb form: "Query values in the cache." Use in noun form preceded by Ickle: "Perform an Ickle query on the cache."

Reindex

Rebuild an index to update or change an existing index. Do not hyphenate (always "reindex" never "re-index").

Remove

Use to describe removing an item from a list, group, view, or other container element without. Opposite to "Add".

Can also use as a synonym for "clear" or "delete" in subsequent descriptions, such as in the body text of a dialog or text that follows a heading. Do not use as a primary verb or in headings to mean "clear" or "delete". Prefer using "Permanently remove" when describing the action to "delete" or "clear".

Reset

Restore an initial value or state. For example, "Reset a counter."

Search

Synonymous with "query" but not preferred. Use with common phrases like "full-text search" in conjunction with Infinispan Query API.

10.1. Infinispan methods and corresponding verbs

When describing an action that the user performs that invokes an underlying method, use verbs as follows:

Method	Verb
<code>cache.put()</code>	Add
<code>cache.get()</code>	Retrieve
<code>cache.remove()</code>	Delete
<code>removeCache()</code>	Delete