

AnswerHub

Social Q&A
for the Enterprise



of the Top 10
StackExchange 1.0 Sites
Now Run on AnswerHub



[Discover Why Now!](#)



CONTENTS INCLUDE:

- › Patterns, Practices and Tools
- › The Patterns
- › Core Patterns
- › Workspace Patterns
- › Code Line Patterns
- › General Guidelines...and more!

Software Configuration Management Patterns

By Steve Berczuk

ABOUT SCM PATTERNS

Software configuration management enables team members to work together more effectively. SCM touches on all aspects of the development process; so well-done SCM can improve productivity. Done poorly, SCM can slow a project down and cause frustration. Effective SCM makes use of more than just the source code management system. It requires that you think about the build and testing process, and also that you continually evaluate how modular your architecture is.

The software configuration management patterns are most applicable to small teams that favor an agile software development approach, but they can help any team identify bottlenecks and work more effectively. If your team isn't agile, but wants to be, following these patterns will provide a framework for your team to develop more agile technical practices.

This Refcard describes some patterns that enable teams to work effectively. These patterns are described in detail in the book *Software Configuration Management Patterns: Effective Teamwork, Practical Integration*.

PATTERNS, PRACTICES, AND TOOLS

The patterns in this Refcard are mostly tool-agnostic. Some tools support the practices implicitly, but any team can implement these patterns using any tool set. You want to think about how you work first, and then look into tools.

The tools you will want to have in place are:

- A source code management system. Either a centralized SCM such as Subversion or a distributed one such as Git will work.
- A build tool that defines dependencies between components and allows for simple build and test execution. Maven and Make are two examples.
- A continuous integration server, e.g., Cruise Control, Anthill, or Bamboo.
- A documentation system such as a Wiki or a well-known location in the SCM repository will do.
- An artifact repository to keep built and third party artifacts. A Maven repository manager or an SCM repository that can store binary files will work.

The operations you need to be able to do are:

- Add code to and checkout code from an SCM repository.
- Commit changes to the SCM repository.
- Create a source code branch.
- Create a checkpoint (a label tag) in the SCM repository.
- Monitor a code line for changes to trigger a build using the CI server.
- Build a project when changes are detected using the CI server.
- Alert the team when a build or its tests fail.

There are a number of commercial and open-source solutions that provide this functionality. If you have an existing tool set it is likely that you can implement the patterns with your tools.

THE PATTERNS

Pattern	Supported By
Main Line	Active Development Line
Active Development Line	Private Workspace Private Versions Task Branch Release Line Release Prep Code Line Code Line Policy
Private Workspace	Third Party Code Line Repository Integration Build Private System Build
Repository	Third Party Code Line
Private Build	Smoke Test
Integration Build	Task Level Commit Smoke Test
Third Party Code Line	
Task Level Commit	
Code Line Policy	
Smoke Test	Unit Test Regression Test
Unit Test	
Regression Test	
Private Versions	Code Line Policy
Release Line	Code Line Policy
Release Prep Code Line	Code Line Policy
Task Branch	Code Line Policy

 Social Q&A
for the Enterprise

 of the Top 10
StackExchange 1.0 Sites
Now Run on AnswerHub

Discover Why Now!

The SCM Patterns can help any team to be more productive, and they especially work well with agile technical practices such as unit testing and continuous integration.

The patterns can be grouped into 3 parts:

- Core patterns that the rest of the language support
- Workspace patterns that describe how developers work
- Code line patterns that describe how to structure supporting code lines to enable delivery of new software from the main line

CORE PATTERNS

The first decision you need to make is how to structure your code lines. The SCM Patterns describe how to work with fewer code lines, emphasizing integration over isolation.

Main Line

The more code lines you have, the harder it is to understand the state of the project. While tools can help manage multiple code lines, the simplest way to minimize the overhead of branch management and context switching is to develop on a single Main Line.

The Main Line pattern can help you to deliver rapidly with a team focused on a single project. With a main line, all changes end up on a single stream of development. This provides the following advantages:

- Reduced merging and synchronization effort; fewer code lines mean fewer merges.
- More consistent integration, reducing the schedule risk of integrating late in the release cycle.

If you are using the Main Line development model, strive to reduce branching to special situations such as:

- **Releases.** Create a release line to manage fixes to released code.
- **Long-lived parallel efforts.** Use a task branch for this.
- **Integration and stabilization.** A release prep code line has advantages over a code freeze for stabilization.

Since each branch is a potential distraction, limit branching to situations where the advantages outweigh potential issues.

The risk with main line development is that developers might commit changes that break the code line. For development on a single code line to be effective you need to give this code line a place where developers can feel confident that the code is working. An active development line provides a framework for a more stable main line.

Active Development Line

An active development line is a main line in which developers use practices to ensure that the code is in a working state. An active development line can support you when you need to do frequent releases and it is essential for agile software development. The patterns that support these practices are:

- Code line policy
- Unit test
- Integration build
- Private workspaces
- Private build

To maintain the activity on the code line, define a code line policy that will keep serious defects out of the code line and integration build, but be willing to ignore trivial defects, as being too strict can paralyze the team.

The Code Line Policy for an Active Development Line could include the following rules:

- The code should run through the Private Build in a Private Workspace successfully before a commit.

- All changes will be accompanied by appropriate tests, or a comment explaining why a test is not being done.
- If the common Integration Build fails, the team will address it immediately.

Since the goal of Active Development line is to err on the side of progress, an occasional broken build is less disruptive than merging uncommitted changes.

WORKSPACE PATTERNS

Build and test practices are essential to maintaining active code lines. The workspace patterns define these.

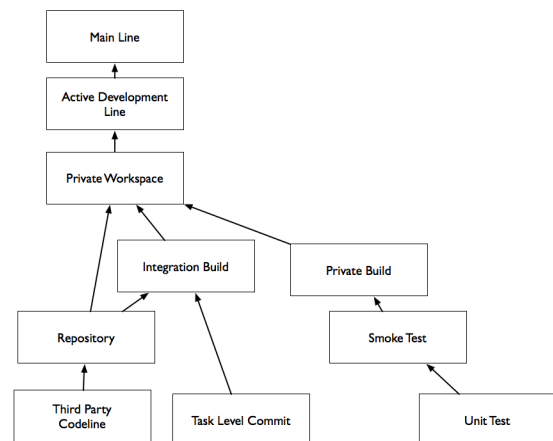


Figure 1: Workspace Patterns

Private Workspace

To maintain an active development line, you need an environment where developers can identify integration issues before code is shared with the team. Developers need control over the state of the code they are working with so that they can work without distraction.

A private workspace is an environment in which developers can build and test before accepting changes from the active development line, or publishing them to the active development line.

A private workspace has all of the dependencies a developer needs to work independently including:

- Code
- Tools
- The correct version of build tools
- The correct version of dependencies
- Configuration files

A developer needs to easily create the workspace from a simple set of instructions using the repository.

The private workspace is useful for developers to get work done. For the team to work effectively, developers need to follow the process in the private build pattern by updating their workspace frequently and committing changes frequently when the code is working. This ensures that the status of the code accurately reflects the state of the active development line.

Repository

To set up a new private workspace or an Integration build, you need to populate it from a repository that contains everything you need to build the code, including:

- Source code
- Build scripts
- Configuration
- Third party components.

The repository can be composed of a number of tools. Source code can be in a source code management system. Components can be in an artifact repository such as a Maven repository or a source code management system.

Ideally, a developer should be able to create a workspace for a project in two steps:

- Check out a copy of the code from your SCM system.
- Build the project.

The only documentation you should need is:

- The path to the project in the SCM System.
- The build command.
- Configuration changes to make for different environments (optional).

You can document the workspace creation process in a well know location in the SCM repository, or on a CMS such as a wiki. A common convention would be to have a Getting Started page. Fewer, self running scripts are better than a longer documented process, but the key attribute of a successful Getting Started process is that it can be executed without assistance from anyone else.

Having all dependencies in a single repository and simple procedure for creating a workspace will minimize the risk of introducing bugs that are related to environmental differences and improve efficiency when people join or move between projects.

Private Build

To avoid breaking the active development line, perform a private build in your private workspace before committing changes. This will allow you to detect integration errors before they affect other developers.

The private build:

- Builds the code.
- Runs smoke tests.
- Runs unit tests.
- Creates a deployable artifact.

The private build should be identical to the integration build, or at least as close as possible. If the integration build skips some tests in the interest of speed, periodically run these tests in the private build.

To avoid checking in code that will break the integration build developers run the Private Build as they develop. Before anything, commit developers should:

- Update their private workspace from the active development line.
- Run the private build.
- Commit their changes only when the build passes.

The private build should be able to grab all dependencies automatically, and not rely on manual installs. A common mechanism for this is to pull dependencies from a Maven or Ivy repository.

Integration Build

Building in a private workspace provides some assurance that all of the code works together. However, you still want an automated mechanism to verify that the code that is in the version management system always builds and passes tests. An integration build runs automatically when changes are detected in the code line.

The integration build:

- Updates the source in an integration workspace.
- Builds the code.
- Runs unit, smoke, and integration tests.

The integration build should be automated, fairly quick, and failures should be addressed immediately. If running a complete suite of tests takes too long, split the integration build into two phases: one that runs smoke tests, and one that runs more thorough unit and regression tests.

Third Party Code Line

All of your locally developed code is in your repository. Code from outside the organization that you depend on should also be there, as you need a way to manage dependencies. For binary dependencies, you can identify versions in your build configurations and use a repository manager. When you need to make customizations to open source code, you might want to manage the source code in your repository. A third party code line is a way you can easily manage local customizations to code:

- Add the third party source to your SCM repository.
- Label the original source.
- Create a branch for your local changes.
- When there is a new release of the third party code, add it to the mainline. Create a new branch for this code.
- Merge any relevant changes from the old branch to the new branch.

Once this is done, create an integration build for the code and a mechanism for developers to reference the third-party artifacts.

Task Level Commit

To help ensure that the integration build line reflects the current state of the code, organize code changes by task-oriented units of work by committing frequently and by associating each task level commit with an issue from your issue tracking system. A task level commit is:

- Small. Commit changes when you have completed a unit of work.
- Frequent. Commit code as often as possible while maintaining working code.
- Associated with a feature being developed. For example; each commit could have an issue number mentioned.

For instance, you might commit after each of these steps:

- Adding a method and unit test.
- Using the new method.

Many issue-tracking systems can associate commits with the issue identifiers either by metadata or by finding issue IDs in the commit comments. Associating each commit with an issue is important in:

- Identifying code changes that went into implementing an issue. This is useful for auditing and research.
- Identifying the effort required for features.
- Helping developers focus their efforts on useful features.

Be sure to update and build code before committing changes to the main line.

Smoke Test

An Integration Build and Private Build use testing to help ensure that your code line is an Active Development Line. To verify that the code line still works after a change, run a Smoke Test after each change as part of the build. A Smoke Test is:

- Quick Running
- Self Scoring
- Providing broad coverage
- Runnable by developers as part of a build-time test

Smoke Tests do not replace all manual quality assurance efforts, but allow for a way to catch common, critical errors quickly after each change.

Unit Test

Smoke Tests provide a quick way to make sure that the application works at a high level. You can rely on smoke tests only if you also have a mechanism to verify that your modules still works after you make a change. Unit Tests are tests that test low level APIs and contracts.

Unit Tests are:

- **Automated and self-evaluating**
- **Fine grained**
- **Isolated, A unit test does not interact with other tests**

Unit tests test the contract that a class has with other components. Run Unit tests while you are coding, before you check in changes, and as part of the build.

Writing unit tests as you code will also help you to identify coupling between modules so that you can remove it if it is inappropriate. Applying practices such as Test Driven Development where you write tests before you write code can be one way to ensure that you have good test coverage.

Unit Tests can also help to identify when integrating conflicting changes from a code merge is successful. If the existing tests and the tests that went with the change you are merging both pass, you can be more confident that you merged changes correctly.

Using a framework like xUnit can simplify your unit testing process.

Regression Test

Unit Tests and Smoke Tests are designed to be fast and meant to be run frequently. You still need a more comprehensive way to ensure that existing code does not get worse as you make other improvements.

Regression tests are a kind of integration test. Regression tests are often driven by problems that you found reactively, and might take longer to run than a build time test should. Ideally, regression tests will be automated.

Regression tests should cover:

- **Problems you find in the QA process**
- **User-reported problems**
- **System level requirements**

When you find an error in a released build, it's a good practice to add a test that identifies the issue to the build.

If the Regression Tests don't take too long to run, add them to the main integration build. Otherwise run them as a second stage build, and consider adding "run regression tests on build" to the code line policy of a Release Line.

CODE LINE PATTERNS

While a single main line that always builds is ideal, there are times when you may want to create branches for certain classes of work to make it easier to keep the main line stable and active. The code line patterns describe these code lines as well as the concept of a code line policy.

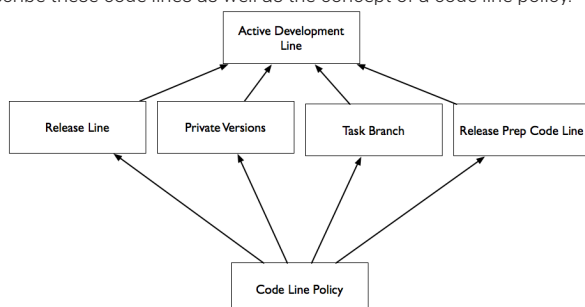


Figure 2: Code Line Patterns

Release Line

Active Development Line development is a simple and powerful approach to efficiently developing software. But developing software in a way that you can deliver at any point in time is challenging. It is useful to have the ability to maintain released versions without interfering with your current development by creating a Release Line. A release line is probably of the more commonly used patterns.

As you release code to customers, identify the version of the release by a tag, and create a branch at that tag when you need to deliver a fix.

When appropriate, integrate changes from the Release Line into the Main Line, either by a merge, or, if the code has diverged significantly, by a parallel change.

A Code Line Policy for a Release Line might include:

- **Each change should be in response to a documented issue.**
- **Each commit must reference an issue identifier.**
- **Changes should be reviewed with another team member before commits.**
- **Any change should be accompanied by a change to an existing test, a new test, or a reason why there is no test.**
- **Before committing a change, a developer should run additional tests such as regression tests, in their Private Workspace.**
- **If the Integration Build fails, every team member stops to help address the issue.**

The details are team dependent; the goal is to have a stable Release Line.

In some ways, development on a Release Line is similar to a waterfall like process, and as such release lines should be used with care.

Release Prep Code Line

The goal of Active Development Line development is to release products from the tip of the code line. In some cases, teams still need extra time to stabilize code before a release while also enabling new work. The traditional approach to this, a code freeze, has a number of down sides, including the opportunity for idle time and/or teams doing work without committing changes to the Main Line.

You can provide for a stabilization period and avoid the downsides of a code freeze by doing the stabilization work on a Release Prep Code line. Create a branch when the majority of the team is ready to start work on the next release, and the current release is feature complete.

If you have a good set of automated tests in place, you will not need to use this pattern. But it is an alternative to a code freeze for teams that still require an extended integration test period.

A Release Prep Code Line is a middle ground between an Active Development Line and a Release Line, so it should have a Code Line Policy that emphasizes that changes should be small and tactical. For example, that policy for Release Prep Code Line might not require the same degree of pre-commit validation as a Release Line commit, but a similar degree of validation. An example policy could be:

- **All Changes should be in response to a bug reported against an issue scheduled for release.**
- **Each commit should identify the issue number it addresses.**
- **Before code is committed, it should be buddy reviewed.**
- **Follow the same pre-commit build and test process as for an Active Development Line.**

Release Prep Code Lines are a stop gap until your team is at a point where it can release code without a code freeze or a long stabilization cycle.

Task Branch

While there are advantages to developing on an Active Development Line, there are times when it's useful to be able to do work in parallel, isolated from the rest of main development work. For these cases, a task branch makes sense. Use a task branch when

- **A subset of your team needs to collaborate on a speculative long-lived task that is a divergence from the main line.**
- **When you are ready to start work on a feature for the next release before the current release is done.**

Because task branches delay integration, use them rarely and only when the benefit outweighs the overhead of the branch. When using a task branch, merge changes from the Main Line frequently so that you are aware of potential conflicts. A Task Branch ends in one of the following ways:

- It's abandoned
- It is merged with the main line. When there is only 1 task branch, the mainline accepts all of the changes.

When working on a Task Branch, it is important to merge changes from the mainline into the task branch frequently. At the end of the task branch, the changes are merged into the main line.

Private Versions

A source code management system's primary purpose is to facilitate collaboration among team members. The facilities it provides to checkpoint steps along the way to implementing a feature make it easy to recover from a mistake and help team members more willing to try things. Since you don't want to check in changes to Active Development Line before code is in a consistent working state, you want to be able to experiment with a complex change locally while still be able to take advantage of features of a version control system.

You can implement Private Versioning either by creating a private branch in the team repository, using a private repository, or by taking advantage of the local history feature of an IDE.

Code Line Policy

At the core of the SCM Patterns are practices that developers follow when working on code from a code line. To make developers aware of rules for a code line, create a Code Line policy to help developers decide what procedures to follow before committing changes to a code line. If possible, automate enforcement of these policies. The code line policy identifies differences between code lines. A code line policy specifies:

- The reason for the code line (for example, Active Development, Fixes for released code, a Task Branch)
- Rules to follow before committing changes. (test, code review, etc)
- Whether the Code Line is long-lived or transient.
- Access restrictions for various roles/individuals/groups

For example:

- An Active Development Line might have a policy that requires Issue Numbers for every commit, Smoke and Unit Tests be run before commit.
- A Release Line might also require that changes be reviewed before being committed, and the all commits should have an associated automated test.

You can enforce most of these rules by tools such as build time steps or SCM triggers. The policies define the agreement between team members about how the code line works.

GENERAL GUIDELINES

This section describes some guidelines to consider when implementing the patterns. The SCM Patterns are a central part of your development process, and can highlight gaps in your process.

Principles

When using the SCM patterns, keep the following principles in mind:

- Fewer Code lines help you to focus on delivering customer value.
- Testing is as important to successful release management as version control.
- Integrate early and often to identify potential problems as early as possible and minimize schedule risk and waster effort.
- When you want to create a branch, consider whether or not there is a simpler solution.

Caveats

A single Active Development Line is the best way to achieve the goals of frequent delivery. The Release Prep Code Line and Task branch are adaptive patterns that provide a way for a team to make progress and help to maintain an Active Development Line, but are workarounds to problems you may have with maintaining stability. Use these patterns sparingly with an eye towards being able to deliver from a single code line.

Testing

Testing is essential to having the SCM Patterns work. Tests validate that the changes in a developer's workspace will not disrupt other developers unexpectedly when they are committed to the Main Line, and also provide a second level of assurance that the integration build works as expected. Though not normally considered part of Software Configuration Management, tests are the best way to validate that the "software configuration" on the code line is correct.

Buddy Reviews

Rather than formal code reviews, teams can benefit from a buddy review process in which a developer asks another developer (the "check-in-buddy") to listen while explains the code changes they are about to check in. While the check-in buddy can sometimes provide suggestions, the main benefit of this process is that the process of having to explain what you did can force critical thinking and help you to identify problems.

Dev Ops and Continuous Delivery

The SCM Patterns can support Dev Ops practices by encouraging Private Build and test environments that transition cleanly to production-like environments.

Balance

The goal of a team using the Active Development Line model is to keep the code line working. Teams should take reasonable care to ensure that developers don't check in broken code. But mistakes happen, and the patterns provide for layers or checking. A developer might forget to update code before building and testing, and an error may make it into the integration build. The goal of the patterns is early detection of errors, not absolute prevention of them. Erring on the side of avoiding errors can also stop progress.

RESOURCES

The SCM Patterns describe approaches to development that are mostly independent of tools, although the right tools will make the process simpler. The approach also relies on a good understanding of techniques such as unit testing and continuous integration.

Continuous Integration Practices

- Refcard 84: Continuous Integration Practices: <http://refcardz.dzone.com/refcardz/continuous-integration>
- Refcard 87: Continuous Integration Servers and tools: <http://refcardz.dzone.com/refcardz/continuous-integration-servers>

Continuous Integration Tools

- Bamboo: <http://www.atlassian.com/software/bamboo/>
- AntHill: <http://www.urbancode.com/html/products/anthillpro/>
- Cruise Control: <http://cruisecontrol.sourceforge.net>
- Hudson: <http://hudson-ci.org>

Build Tools

- Maven: <http://maven.apache.org>
- Ant: <http://ant.apache.org>
- Buildr: <http://buildr.apache.org>
- Gradle: <http://www.gradle.org>

Development Practices

- Test Driven Development Resources: <http://www.testdriven.com>
- Dev Ops Resources: <http://devops.com>
- Continuous Delivery: <http://continuousdelivery.com>

Unit Testing

- Refcard 28: Java Unit Testing tools: <http://refcardz.dzone.com/refcardz/junit-and-easymock>
- The JUnit web site has links to other Unit Testing frameworks: <http://www.junit.org>
- Guidelines on Unit Testing: <http://www.extremeprogramming.org/rules/unittests.html>
- Guidelines on Test Driven Development: <http://www.testdriven.com>

SCM Tools

- Refcard 94: Getting Started with Git: <http://refcardz.dzone.com/refcardz/getting-started-git>
- Subversion: <http://subversion.apache.org>

CONCLUSION

Maintaining a single working code line sounds simple in theory, but has many practical challenges. The advantages, such as an approach offers in terms of simplicity and minimizing integration costs across the team, can make the cost worthwhile. Doing Software Configuration Management effectively involves more than just SCM tools and code line management, but also your approach to development, including testing.

ABOUT THE AUTHORS



Steve Berczuk is the author of *Software Configuration Management Patterns: Effective Teamwork, Practical Integration* and is a recognized expert in software configuration management and agile software development. Steve is passionate about helping teams work effectively to produce quality software. Steve is an engineer at Humedica where he's helping to build next-generation SaaS-based clinical informatics applications. He has an M.S. in Operations Research from Stanford University and an

S.B. in Electrical Engineering from MIT, and is a Certified Practicing Scrum Master. Contact Steve at steve@berczuk.com or visit www.berczuk.com, or read his blog at blog.berczuk.com.

RECOMMENDED BOOK



This book describes many of the common problems organizations face when using SCM in a way that is not consistent with the needs of their organization. The book presents the patterns organized in a pattern language that shows how the various structures that you need to use for an effective SCM system build on each other.

[BUY NOW](#)

Browse our collection of over 150 Free Cheat Sheets

Free PDF

Upcoming Refcardz

Apache HTTPD
Mongo DB
Essential Camel Components
Debugging Patterns



DZone communities deliver over 6 million pages each month to more than 3.3 million software developers, architects and decision makers. DZone offers something for everyone, including news, tutorials, cheat sheets, blogs, feature articles, source code and more.

"DZone is a developer's dream," says PC Magazine.

DZone, Inc.
150 Preston Executive Dr.
Suite 201
Cary, NC 27513

888.678.0399

919.678.0300

Refcardz Feedback Welcome

refcardz@dzone.com

Sponsorship Opportunities

sales@dzone.com

ISBN-13: 978-1-936502-64-6
ISBN-10: 1-936502-64-X



\$7.95