

## Lab Exercise

# Git: A distributed version control system

### Goal

This lab is intended to demonstrate basic usage of `git`, a distributed version control system. You are supposed to work in pairs on one computer, simulating a small software team developing a simple program component under version control. The lab also demonstrates the XP practices of Collective Ownership and Continuous Integration. The lab is divided into two parts. In part one you will learn git working from the command line. In part two you will instead do all version control directly from within Eclipse.

### Note!

- During the lab, keep notes of what you do.
- You may use a text file instead of paper for your notes, if you think that is easier.
- Write down the commands that you use and brief descriptions of the results.
- Write down brief answers to all questions appearing in this lab text.
- Be prepared to discuss your notes with the lab coach.
- As part of the explanation we sometimes refer to CVS and compare git commands with CVS commands. This is because you might already be familiar with CVS and the comparison will make it easier for you to understand. If you are not familiar with CVS, just ignore this comparisons and focus on the git commands.

## 1 Working from the command line

Later in this lab and also in the project, you will use version control tools in Eclipse. However, in the first part of this lab exercise you will work with the version control tools directly from the command line, and edit files using an ordinary text editor. This will give you a better insight into how the tools work, than if you had only used them via Eclipse. If you are unfamiliar with unix concepts and commands, you can find information in "Introduktion till LTH:s Unixdatorer" (<http://www.ddg.lth.se/perf/unix/unix-x.pdf>), and you can also google different commands.

## 2 Introduction

In this lab you will meet Alice and Bob, i.e. you will play the roles of Alice and Bob to try out Git. Alice and Bob are planning to move their development to the new popular version control system Git. They want to start working in a simple way, similar to CVS, with a common central repository that they both can commit changes to, or rather, using Git parlance, they will *push* changes to it. Since Git works a bit differently from CVS, they draw two figures to get a better understanding of similarities and differences. In figure 1 it is shown how `checkout` is used in CVS to create a local workspace from the

common repository. In Git, the command `clone` is used instead, which creates both a local repository and an associated workspace. The local repository is a clone of the common repository, and contains all the versions of the software. The common repository is *bare*, meaning that it does not have any associated workspace.

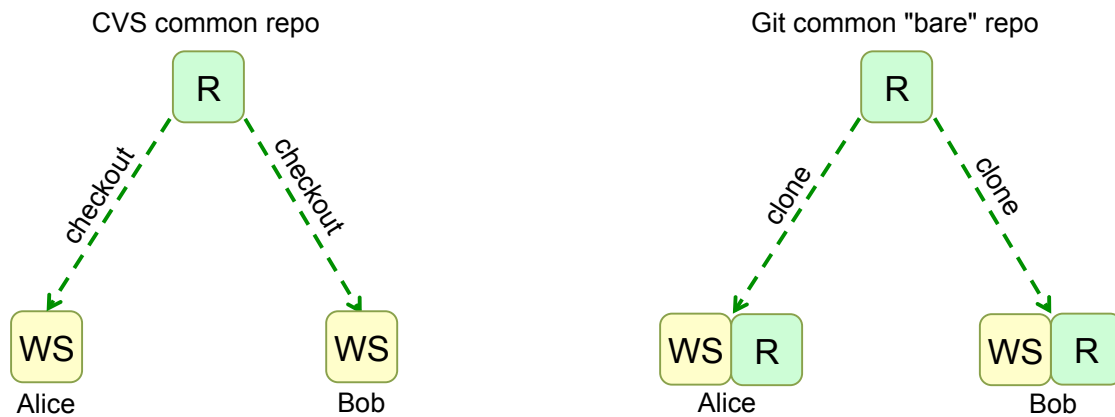


Figure 1: Creating a local workspace. In Git, a local repository is created along with the workspace.

Figure 2 shows how a developer works. In CVS, the developer *updates* the workspace to get the latest version from the repository, edits files, and then *commits* the changes to the common repository. In Git, the developer *pulls* the latest version from the common repository to his/her own local repository/workspace, then edits files, then *commits* the changes to the local repository, and finally *pushes* changes in the local repository back to the common repository.

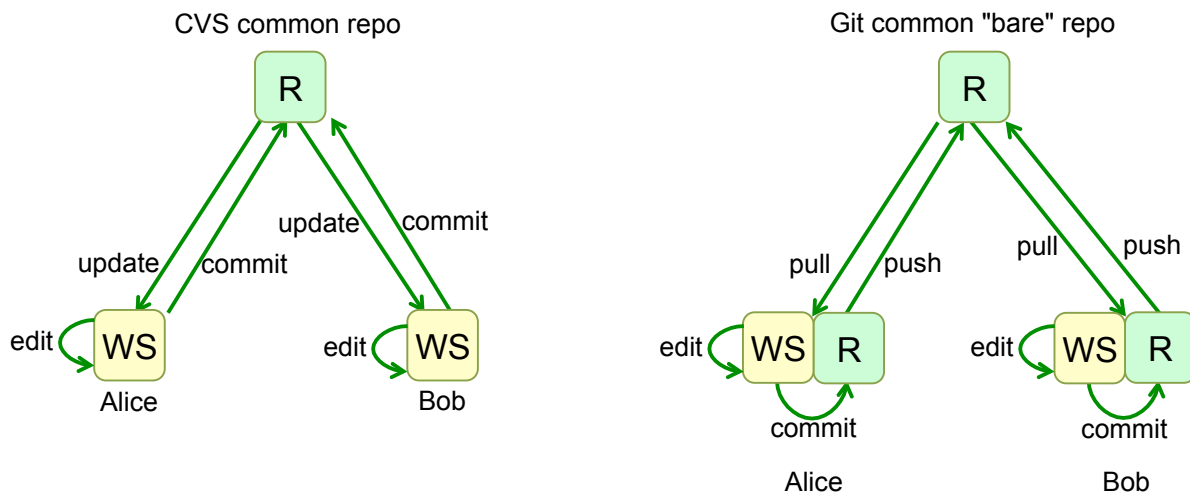


Figure 2: Working with CVS and Git. In Git, the changes are first committed to the local repository. Then the changes are pushed to the common repository.

Now, study the two figures and make sure you understand what *clone*, *pull*, *commit*, *push*, and *bare* means in Git. Write down your own descriptions here:

<b>clone</b>	
<b>pull</b>	
<b>commit</b>	
<b>push</b>	
<b>bare</b>	

Alice and Bob now have some doubts about Git because it would seem that it could take a lot of space to have the complete repository locally. And in CVS you can check out a single module from the repository, whereas in Git you get a workspace for the latest version of the complete repository. However, after consulting with others, they find out the following:

- Git uses compression algorithms to keep down the size of the repositories. But to keep down the size it is also important to commit only source data, like source files and text files, and avoid committing large generated binary files, like class files and jar files.
- In CVS it is common to have a large repository with many submodules, e.g., for different products or packages. But in Git, you typically have many smaller repositories instead.

Alice and Bob still think that Git looks a bit more complicated than CVS, so they wonder about what the advantages could be. Again, they consult with others and get a lot of different answers, for example the following:

- You can work locally, committing to your own local repo, without any network connection.
- You can create branches locally, without needing to make them visible to others.
- You can create your own repositories locally, so you can get the benefits of version control for your own private work, not just for collaboration.
- Instead of using a central repository, you can push and pull directly between developer repositories, creating your own work policies.
- If the central server breaks down, you will still have all your data, and can keep on working.
- It's the most popular version control system right now, so you get a lot of cool support from source code hosting providers like GitHub, BitBucket, Google Code, and others.

Encouraged by this, Alice and Bob think it is worth giving Git a try. They will build a simple HelloWorld application and version control it using Git. In the following, you should play the roles of Alice and Bob, as indicated by the box in the margin.

### 3 Alice starts working locally

Alice is eager to get started. She rushes ahead and creates a new directory for the project in her pvglab2 directory, and creates an empty README file for the project.

Alice

```
$ mkdir pvglab2
$ cd pvglab2
$ mkdir HelloProject
$ cd HelloProject
$ touch README
```

Alice thinks about this wonderful news that she can have her own local git repository, and get version control on her own code. She decides to try this out before contacting Bob about any collaboration.

### 3.1 Configuring git

Alice has never worked with git before, so she first configures git with her name, email, and the default text editor she would like to use (`nano` is still her favorite text editor):

Alice

```
$ git config --global user.name "Alice Wonderland"
$ git config --global user.email alicedat12@student.lth.se
$ git config --global core.editor nano
```

She then checks her current configuration, to make sure the new info is there:

```
$ git config --list
```

She wonders where this information is stored, and finds out that it is in a file called `.gitconfig`, in her home directory, so she looks at its content:

```
$ less ~/.gitconfig
```

### 3.2 Creating a local repository

Alice creates a local repository simply by doing the git command `init` in her new project directory, adding all files (currently just the README file), and committing them to the new repository:

Alice

```
$ cd HelloProject
$ git init
$ git add .
$ git commit -m "New project"
```

She wonders where git placed the repository, and finds it in a subdirectory `.git`. She takes a look at the contents of the repo, just out of curiosity:

```
$ ls -l .git
```

She notes that there are files and directories called `HEAD`, `branches`, `config`, etc., and she thinks this looks just like the kind of stuff you would expect in a version control repository.

### 3.3 Working with the local repository

Alice now starts working with her local repository. She edits the README file, and commits again, checking the git status in between the commands, and checking the commit log:

Alice

```
$ git status
$ nano README
$ git status
$ git add README
$ git status
$ git commit -m "Updated README"
$ git status
$ git log
```

She notes some things:

- She is working on a branch called *master*
- Before committing she needs to add new and changed files to the so called *staging area*
- Changed files can be staged by using the `-a` option for *commit*, instead of explicitly doing an *add*
- When there is nothing to commit, the working directory is said to be *clean*

Alice also finds out that each commit is identified by a long hexadecimal string, like

**6b8736fe84cb3e842352af055e0b3948e602fc80**

She finds out that this is a hash value of the contents of the committed version and its complete history of earlier versions. Alice also finds out that these hash values (produced by the SHA-1 secure hash algorithm) are sufficiently unique, so that it is extremely unlikely that two different commits would get the same hash code. In fact, she read somewhere that a hash collision might occur if you had  $2^{80}$  commits, which incidentally is the same order of magnitude as the number of atoms in the universe. So Alice feels pretty confident it will not happen in her project. She also finds out that Git uses these hash values internally to identify files, directories, etc., to be able to decide extremely quickly if two files/trees/commits, etc. are equal or not, i.e., without having to actually compare the contents.

Alice notes that commands like *commit* and *push* use seven-digit hexadecimal numbers, like **6b8736f**, and she finds out that this is simply a way of abbreviating the long hash values by only showing the first seven digits.

## 4 Setting up a common repository

Alice would now like to create a common repository, so that she can invite Bob to start collaborating with her on the new project. She will simply place the common repository in Bitbucket cloud. She will create the repository on her user account there, and give Bob write access. This way, Bob can push to the repository, even if it is located on Alice's account.

When logged in to your Bitbucket account at <http://bitbucket.org>, you will find online help and tutorials that you can look at if needed.

### 4.1 Creating the common repository

Alice now log in to Bitbucket and creates a new repository. She calls the new repository **HelloProject**. She makes it private. She also, already now, gives Bob write access to her repo (in "Settings" and "User and group access") so that he will be able to push into it.

Alice

### 4.2 Connecting the developer repository to the common repository

Alice now needs to connect her developer repository to the new common repository and push its contents to it. She goes to the developer repository and performs the following magic for doing this:

Alice

```
$ cd pvglab2/HelloProject
$ git remote add origin https://<username>@bitbucket.org/<username>/helloproject.git
$ git push -u origin master
```

where `<username>` is Alice's Bitbucket username.

These commands will make it look like the developer repository was cloned from the common repository (although we in fact created the repositories in the opposite order).

Out of curiosity, Alice looks up the more precise meaning of these commands and finds out the following:

- The *remote* command sets up the common repository as the *origin* of the developer repository.
- The *push* command pushes the *master* branch of the developer repository into the origin repository.
- The *-u* option configures her master branch to *track* the master branch in the common repository. This has the effect that when she in the future does a **git push**, the changes will end up in the right place in the common repository.

Alice does not entirely grasp all this, but is happy with accepting that the magic probably does what it should. To check that everything works, Alice does a small change to the README file, commits it, and pushes it to the common repository.

```
$ nano README
$ git commit -a -m "Updated README"
$ git push
```

She also tries the pull command.

```
$ git pull
```

Of course, nothing new is pulled, since Bob has not started working yet.

### Avoiding writing the password?

Alice thinks it is a bit tedious to have to write the **bitbucket** password for every pull or push. She knows this can be avoided by ssh and generating an ssh key, but she thinks this is too much work for this short experiment with git.

## 5 Alice invites Bob to work on the project

Alice now invites Bob to work on the project. He starts by configuring git in a similar way as Alice did.

**Bob**

### 5.1 Cloning the common repository

Bob clones the common repository into a suitable place in his own account, and checks the contents of his working directory.

**Bob**

```
$ cd pvglab2
$ git clone https://<Bob username>@bitbucket.org/<Alice username>/helloproject.git
$ cd HelloProject
$ ls -a
```

## 5.2 Bob starts working

Bob starts working. He edits the README file some more, commits it, and pushes back the changes to the common repository. Write down the commands he uses here:

Bob

## 5.3 Alice pulls the new changes

Alice pulls down Bob's changes, and takes a look at the commit log.

Alice

```
$ git pull
$ git log
```

## 5.4 Concurrent development: using stash

Alice and Bob now start experimenting with developing at the same time. Just like you should do *update* before you do *commit* in CVS, you should do *pull* before you do *commit* and *push* in Git. And just like in CVS, you should make sure that the code is clean (compiles and tests) before committing and pushing.

First, try the following scenario where Alice and Bob edit the same file:

- Alice edits the README file and saves it.
- Bob edits the README file and saves it.
- Alice pulls, commits, and pushes.
- Bob pulls in order to merge with the latest version before he commits and pushes.

Alice

Bob

Alice

Bob

What happens? It turns out that you cannot pull if your working directory contains uncommitted files that need to be merged with the pulled changes. Bob now has two options:

**commit** he can commit his changes and then pull

**stash** he can *stash* away his changes, then do pull, and later *apply* the stashed changes, merging them with the new version of the files.

The *stash* command takes all the uncommitted changes to the working directory and moves them away to a safe place, the *stash*, so that all edits are undone, and the working directory is set back to the previous *clean* state (with no uncommitted changes). The *stash apply* command takes the changes from the stash, and applies them on the working directory again, i.e., merges them with the files in the working directory.

Bob tries the stash option, checking the contents of the README file after each step. If necessary, he edits the README file before the commit, to make sure it has the desired merged content.

Bob

```

$ git stash
$ more README
$ git pull
$ more README
$ git stash apply
$ more README
$ nano README
$ git commit -a -m "...
$ git push

```

Bob realizes that *stash* is like a light-weight anonymous branch.

## 5.5 Continued concurrent development

Alice and Bob continue with their development, creating the HelloWorld application. They experiment with different ways of getting into merge situations, making sure that Alice also tries out the stash command.

Alice and Bob then reflect on what they have learnt, by writing down descriptions of git concepts and commands:

<i>repository</i>	
<i>working directory</i>	
<i>staged area</i>	
<i>stash</i>	
<i>commit identifier</i>	
git config	
git init	
git clone	
git status	
git log	
git add	
git commit	
git push	
git pull	
git stash	
git stash apply	



## 6 Part II - Using Git from Eclipse

Now you know the basics of git and how to work with Bitbucket cloud as the common repository. In this part of the lab, you will create a Java project in Eclipse and put it under version control using git and Bitbucket. All git commands will be executed from within Eclipse. However, feel free to in parallel see what happens in you local git repository from the command line.

### 6.1 What to do

Study the YouTube clip made by Patrik (<https://www.youtube.com/watch?v=JNiWJhi80Cc>). Do "exactly" the same thing using your bitbucket account, e.g. use "Alice" account and give "Bob" write access. Also, start two instances of Eclipse, one from each of your unix accounts, to make it as real as possible simulating two developers working in parallel on the same project.

Make sure you experience several possible scenarios of parallel modifications and how they can be solved using git and Eclipse.

## 7 Learning more

Alice and Bob now discuss some additional things they would like to learn about Git, for example how to use branches in Git. As a starting point, they take a look at the official Git site: <http://git-scm.com>, then they decide it is time for a break.