

Git

(Why not CVS? ... because Git.)

Karel Zak

<kzak@redhat.com>

Copyright © 2007 Karel Zak.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.

Agenda

1 Intro

- Development model

2 Implementation

- Naming revisions

3 Repositories

4 Branches

5 Real life with Git

6 Commands

- Commits and patches

7 Misc

What is Git?


"I'm an egotistical bastard, and I name all my projects after myself. First 'Linux', now 'git'." (Linus Torvalds)

- fast distributed revision control system
- unusually rich command set
- provides both high-level operations and full access to internals
- originally created by Linus Torvalds for kernel development
- design was inspired by BitKeeper and Monotone
- GNU General Public License, version 2

Basic features

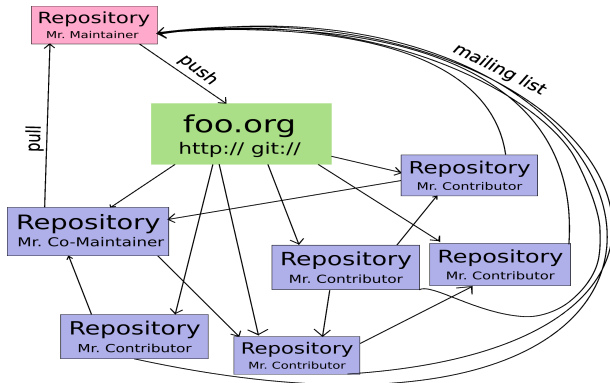
- distributed development model
- support for non-linear development (branching and merging)
- ready for large projects (very good performance)
- repositories can be easily published (git://, ssh://, http://, rsync://, ...)
- cryptographic authentication of history (GPG-signed tags)
- internally objects are addressed by content (SHA-1) – no filenames
- local branches are local only (off-line work)

Centralized model



- extra policy for write access
- SCM is not development tool, but source code archive only
- every change has impact to all developers

Distributed model



- maintainer has full control over primary (his) repository
- support for non-linear development
- repositories can be easily published (git://, ssh://, http://, ...)

Git improves your work manners and habits

- branching and merging is cheap
 - you can prototyping
 - you can collaborate with others developers on incomplete and unstable stuff
 - you can easily (e.g. every day) rebase your changes to new upstream code
 - often merging (rebase) minimize conflicts between your patches and upstream
- small patch is the best patch (patch per feature/change)
 - reviewers hate huge patches
 - well separated feature or change is easy to revert

Internal objects

All objects are content-addressable by SHA-1.


commit refers to “tree” and “parent” (connection into the project history)

tree represents the state of a single directory (list of “blob” objects and subtrees)

blob contains file data without any other structure

references are human readable names for commits; for example tags, see files under `.git/refs/heads/`.

Internal objects



commit – connection between “tree” and “parent”

tree – state of a single directory

blob – contain file data

Trust

- everything is content-addressed and based on SHA-1
- two trees are same when HEAD SHA-1 are same
- tags could be GPG-signed

```
$ git tag -v v2.13
object 49ef7acdf77066ed05a6c828c261d332c4f54644
type commit
tag v2.13
tagger Karel Zak <kzak@redhat.com> Tue Aug 28 01:01:35 2007 +0200
```

```
stable release v2.13
```

```
gpg: Signature made Tue 28 Aug 2007 01:01:35 AM CEST using DSA key ID DC06D885
gpg: Good signature from "Karel Zak <kzak@redhat.com>"
```

Object reference

SHA-1 40-hexdigit object name

tag human readable name for commit

`commitn` N-th parent

`commit~n` N-th generation grand-parent of the named commit object, following only the first parent.

ref@{date} specify the value of the ref at a prior point in time

`:/text` commit whose commit message starts with the specified text

HEAD refers to the head of the current branch

```
rev~3 = rev^^^ = rev1^1^1
```

```
$ git reset HEAD^
```

Ranges

r1..r2 commits reachable from r2 but exclude the ones reachable from r1

r1...r2 set of commits that are reachable from either one of r1 or r2 but not from both

```
$ git log v2.13..v2.14
```

"tree-ish"

Lots of commands take a tree as an argument. A tree can be referred to in many different ways, by:

- name for that tree
- name of a commit that refers to the tree
- name of a branch whose head refers to that tree

Create a repository

- create a new repository

```
$ mkdir project  
$ cd project  
$ git init
```

- clone an existing remote repository ("origin" repository)

```
$ git clone http://foo.com/project
```

- add a next remote repository

```
$ git config remote.bar.url git://bar.com/project  
$ git config remote.bar.fetch master:refs/remotes/bar/master  
$ git fetch bar
```

Visualisation

File Edit View Help

- master: setarch: add long options to setarch and update manpage
- docs: update AUTHORS file
- fdisk: doing useless ioctl when editing an image
- rename: add description about option -U to manpage
- rename: remove useless variable
- build-595: remove files that are no longer delivered from git
- cal: add description about option -U to manpage
- mount: chain of symlinks to fstab causes use of pointer after free
- fstab: fix typo
- fstab: rename: add description about option -U to manpage
- rename: remove useless variable
- mount: chain of symlinks to fstab causes use of pointer after free
- cal: add description about option -U to manpage
- build-595: remove files that are no longer delivered from git
- setarch: adding groff symlinks to setarch manual page
- mount: improve error message when helper program not present
- fdisk: fix typo
- fdock: typo in man page
- build-595: fix locale_dir (unsupported by autoconf < 2.60)
- docs: fix stable branch name in README.devel
- topic: swap-on-suspend: Merge branch 'master' into topic:swap-on-suspend
- setarch: adding groff symlinks to setarch manual page
- mount: improve error message when helper program not present
- docs: fix stable branch name in README.devel
- fdisk: fix typo
- autogen.sh reports versions of autotools now
- build-595: set AC_PREREQ to 2.60, increment version to 2.14
- fdock: typo in man page
- 02-13: build-595: release++
- docs: update ReleaseNotes
- docs: update and sort AUTHORS file

Karel Zak <kzak@redhat.com>	2007-09-04 09:45:10
Karel Zak <kzak@redhat.com>	2007-09-05 11:55:50
Pascal Terjan <pterjan@linusmf.org>	2007-09-04 00:10:16
Li Zefan <liz@cn.fujitsu.com>	2007-09-04 09:33:36
Li Zefan <liz@cn.fujitsu.com>	2007-09-04 09:20:25
Laltonf Jones <laltonf@mmjgroup.com>	2007-08-28 23:56:41
Li Zefan <liz@cn.fujitsu.com>	2007-09-03 08:25:40
Norbert Buchmüller <norbi@gnix.ch>	2007-09-02 22:08:53
Li Zefan <liz@cn.fujitsu.com>	2007-09-04 09:33:36
Li Zefan <liz@cn.fujitsu.com>	2007-09-04 09:20:25
Norbert Buchmüller <norbi@gnix.ch>	2007-09-02 22:08:53
Li Zefan <liz@cn.fujitsu.com>	2007-09-03 08:25:40
Laltonf Jones <laltonf@mmjgroup.com>	2007-08-28 23:56:41
Arkadiusz Miskiewicz <arekm@maven.pl>	2007-08-28 09:16:57
Laltonf Jones <laltonf@mmjgroup.com>	2007-08-09 21:37:22
Karel Zak <kzak@redhat.com>	2007-09-02 15:02:55
A. Corfa <apcorfa@gnix.net>	2007-08-29 15:07:10
Karel Zak <kzak@redhat.com>	2007-09-02 13:00:23
Karel Zak <kzak@redhat.com>	2007-09-02 11:43:53
Karel Zak <kzak@redhat.com>	2007-09-04 15:43:33
Arkadiusz Miskiewicz <arekm@maven.pl>	2007-08-28 09:16:57
Laltonf Jones <laltonf@mmjgroup.com>	2007-08-09 21:37:22
Karel Zak <kzak@redhat.com>	2007-09-02 11:43:53
Karel Zak <kzak@redhat.com>	2007-09-02 15:02:55
Karel Zak <kzak@redhat.com>	2007-09-02 13:43:31
Karel Zak <kzak@redhat.com>	2007-09-02 13:10:06
A. Corfa <apcorfa@gnix.net>	2007-08-29 15:07:10
Karel Zak <kzak@redhat.com>	2007-08-28 00:41:54
Karel Zak <kzak@redhat.com>	2007-08-28 00:38:57
Karel Zak <kzak@redhat.com>	2007-08-28 00:15:49

SHORT ID: e4d8671e228720849a7e050ee1108e163a57b255

Highlight Comments touching paths: OR in view None OR None

Search Patch Tree

Diff Old version New version

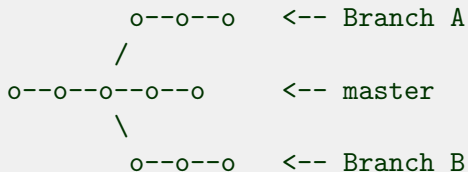
Author: Pascal Terjan <pterjan@linusmf.org> 2007-09-04 00:10:16
 Committer: Karel Zak <kzak@redhat.com> 2007-09-05 11:52:28
 Parent: 24e5745f5594e9545533e528d6f90b7fe00a8 (rename: add description about option -V to manpage)
 Diff: [diff: 11e11288972473361e22c3f7b1e55f1e811941c](#) (docs: update AUTHORS file)
 Branch: master
 Follows: [v2-13](#)
 Precedes:

```
fdisk: doing useless ioctl when editing an image

When editing a disk image, fdisk wants to ask the kernel to reread the
partition table which is useless and produces an error. A wrong exit
```

Comments
fdisk/fdisk.c

Branches



branch is line of development

branch head is a reference to the most recent commit on a branch

Branches, remote-tracking branches, and tags are all references to commits.

Manipulating branches

- **git branch** lists, creates, or deletes branches
- **git checkout <branch>** makes the current branch <branch>, updating the working directory
- **git checkout -b <branch>** creates a new branch <branch>check it out
- **git show-branch** shows branches and their commits
- **git diff <branch>..<branch>** diffs between branches

Rebase branch

Before

```
      A---B---C topic
      /
D---E---F---G master
```

Command

```
$ git rebase master topic
```

After

```
      A---B---C topic
      /
D---E---F---G master
```

Edit 2nd commit from the top

- 1 create a temporary branch (rewind to the commit in question):

```
$ git checkout -f -b tmp HEAD~2
```

- 2 reset old changes and redo the commit

```
$ git reset HEAD~  
$ vim foo.c  
$ git commit -a -c ORIG_HEAD
```

- 3 replay the later changes to the master:

```
$ git rebase --onto tmp master~2 master
```

- 4 clean up (delete) the temporary branch

```
$ git branch -D tmp
```

Changes in project history

- the latest patches – (`git reset`) reset current HEAD
- deep in project history
 - rebase
 - impact to all rebased commits (new SHA-1)
 - impact to all tags
 - useless for publicly pushed changes
 - patch revert (`git revert`)
 - zero impact to project history

Send a patch

Basic rules:

- one patch per e-mail
- don't use stupid e-mail clients (e.g. Outlook)
- (don't use attachments)
- export patches by `git format-patch`
- send patches by `git send-email`
- well formatted patch is possible to apply by `git am`
- don't forget to keep correct authorship (e.g when you are not author of the patch)
- use commit messages – a patch without comment is incomplete crap

Syntax

- `git <commandname> [options]`
- `git-<commandname> [options]`

High level

```
$ git commit -a -s -m "cool change"
```

Low level

```
$ git rev-list --pretty=oneline v2.13..
```

Basic commands

git init creates an empty repository at `./`.git

git add adds file contents to the next commit

git reset resets current HEAD

git status shows the working tree status

git commit records changes to the repository

git log shows commit log

git show shows commit (or another object)

git format-path exports a change

git send-email sends patch(s)

git am applies a series of patches from a mailbox

Commit changes

```
git commit [options] [<file>]
```

- **-a** commits all modified or deleted files
- **-s** adds Signed-off-by line at the end of the commit message
- **-c ORIG_HEAD** reuses a commit message (e.g. from previously reset commit)
- **-author "Random Hacker <rand@hacker.com>"** overrides the author name
- starts **\$EDITOR** for commit message (or **-m "<commit message>"**)

```
$ git commit -a -s
```

Export patches to files

```
git format-patch [options] <since|range>
```

- creates one file per patch
- created patches are usable by **git am**

```
$ git format-patch -o ~/ HEAD~5  
/home/kzak/0001-setterm-opened-file-leaving-unclosed.patch  
/home/kzak/0002-sfdisk-opened-files-leaving-unclosed.patch  
/home/kzak/0003-blockdev-fix-opened-file-leaving-unclosed.patch  
/home/kzak/0004-tailf-opened-file-leaving-unclosed.patch  
/home/kzak/0005-tests-use-losetup-s.patch
```

Patch description

```

From bfdb8be5c49d8fad25118fb4416ab2a68fc3a16 Mon Sep 17 00:00:00 2001
From: Karel Zak <kzak@redhat.com>
Date: Thu, 25 Oct 2007 12:29:51 +0200
Subject: [PATCH] losetup: canonicalize loopfile name

When setting up a loop device, canonicalize the loop file
name. This simplifies a later identification of loop file names
when querying the loop devices.

Co-Author: Matthias Koenig <mkoenig@suse.de>
Signed-off-by: Matthias Koenig <mkoenig@suse.de>
Signed-off-by: Karel Zak <kzak@redhat.com>
---
mount/Makefile.am | 4 +--
mount/lomount.c | 14 ++++++++--
2 files changed, 14 insertions(+), 4 deletions(-)

diff --git a/mount/Makefile.am b/mount/Makefile.am
index 57a5af8..46bcb5a 100644
--- a/mount/Makefile.am
+++ b/mount/Makefile.am
@@ -25,8 +25,8 @@ umount_LDFLAGS = $(SUID_LDFLAGS) $(AM_LDFLAGS)

swapon_SOURCES = swapon.c swap_constants.h $(utils_common)

-losetup_SOURCES = lomount.c sundries.c xmalloc.c loop.h lomount.h xmalloc.h \
+ sundries.h
+losetup_SOURCES = lomount.c sundries.c xmalloc.c realpath.c \
+ loop.h lomount.h xmalloc.h sundries.h realpath.h
losetup_CPPFLAGS = -DMAIN $(AM_CPPFLAGS)

mount_LDADD = $(LDADD_common)
diff --git a/mount/lomount.c b/mount/lomount.c
index cea3aed..b3c16fb 100644
--- a/mount/lomount.c
+++ b/mount/lomount.c
@@ -24,6 +24,7 @@
#include "nls.h"
#include "sundries.h"
#include "xmalloc.h"
+#include "realpath.h"

#define SIZE(a) (sizeof(a)/sizeof(a[0]))

```

header

commit
message

tags

diffstat

patch

Send patches by e-mail

```
git send-email [options] <file|dir>
```

Takes the patches given on the command line and emails them out.

- no attachments
- no broken patch format
- correct subject line

```
$ git send-email --to "God <father@heaven.com>" \  
~/0001-make-this-world-better.patch
```

Browsing changes

git log shows commit logs

git show shows one or more objects (blobs, trees, tags and commits)

git blame shows what revision and author last modified each line of a file


git whatchanged shows logs with difference each commit introduces

```
$ git log v2.5..                # commits since v2.5
$ git log test..master         # commits reachable from master
                                # but not test
$ git log --since="2 weeks ago" # commits from the last 2 weeks
$ git log Makefile             # commits which modify Makefile
$ git log --pretty=format:"%h [%an]" # commit log in format
                                # "sha-1 [Author Name]"
$ git blame -L 10,15 foo.c     # who modified code between lines
                                # 10 and 15
$ git show c1a47c171b         # shows selected object (commit)
```

Gitweb

File Edit View History Bookmarks Tools Help

http://git.kernel.org/?p=utils/util-linux-ng/util-linux-ng.git

/pub/scm / utils/util-linux-ng/util-linux-ng.git / summary 

summary | [shortlog](#) | [log](#) | [commit](#) | [commitdiff](#) | [tree](#)

commit search:

description The util-linux-ng code repository. See <http://kernel.org/~kzak/util-linux-ng/> for more details.

owner Karel Zak

last change Wed, 5 Sep 2007 09:55:50 +0000

URL <git://git.kernel.org/pub/scm/utils/util-linux-ng/util-linux-ng.git>
<http://www.kernel.org/pub/scm/utils/util-linux-ng/util-linux-ng.git>

shortlog

5 days ago	Karel Zak	docs: update AUTHORS file [master]	commit commitdiff tree snapshot
5 days ago	Pascal Terjan	fdisk: doing useless ioctl when editing an image	commit commitdiff tree snapshot
5 days ago	Li Zefan	rename: add description about option -V to manpage	commit commitdiff tree snapshot
5 days ago	Li Zefan	rename: remove useless variable	commit commitdiff tree snapshot
5 days ago	LaMont Jones	build-sys: remove files that are no longer delivered ...	commit commitdiff tree snapshot
5 days ago	Li Zefan	cal: add description about option -V to manpage	commit commitdiff tree snapshot
5 days ago	Norbert Buchmüller	mount: chain of symlinks to fstab causes use of pointer ...	commit commitdiff tree snapshot
7 days ago	Arkadiusz Miskiewicz	setarch: adding groff symlinks to setarch manual page	commit commitdiff tree snapshot
7 days ago	LaMont Jones	mount: improve error message when helper program not ...	commit commitdiff tree snapshot
7 days ago	Karel Zak	docs: fix stable branche name in README.devel	commit commitdiff tree snapshot
8 days ago	Karel Zak	fdisk: fix typo	commit commitdiff tree snapshot
8 days ago	Karel Zak	build-sys: autogen.sh reports versions of autotools now	commit commitdiff tree snapshot
8 days ago	Karel Zak	build-sys: set AC_PREREQ to 2.60, increment version ...	commit commitdiff tree snapshot
8 days ago	A. Costa	flock: typo in man page	commit commitdiff tree snapshot
13 days ago	Karel Zak	build-sys: release++ v2.13	commit commitdiff tree snapshot
13 days ago	Karel Zak	docs: update ReleaseNotes	commit commitdiff tree snapshot

Done Adblock

References



Git User's Manual

<http://www.kernel.org/pub/software/scm/git/docs/user-manual.html>



A tutorial introduction to git

<http://www.kernel.org/pub/software/scm/git/docs/tutorial.html>



The perfect patch

<http://www.zip.com.au/~akpm/linux/patches/stuff/tpp.txt>

The end.

Thanks for listening.