



Git Inside

s3lph

s3lph@s3lph.me

@s3lph@chaos.social

June 16, 2019

CoSin 2019

What?

- > .git Directory
- > References
- > Objects
- > Push & Fetch

Why?

1. Want to know how Git works
2. Read lots of documentation
3. Write your own Git repo parser
4. ???
5. This talk

Porcelain Commands – Overview

- > git init
- > git add
- > git commit
- > git branch
- > git checkout
- > git reset
- > git rm
- > git status
- > git fetch
- > git pull
- > git push
- > git branch
- > git cherry-pick
- > git merge
- > git rebase
- > ...

.git Directory Overview

```
.git/  
├── branches/  
├── config  
├── HEAD  
├── hooks/  
├── index  
├── info/  
│   └── logs/  
├── objects/  
└── refs/
```

.git Directory Overview

```
.git/  
├── branches/  
│   └── origin
```

.git Directory Overview

```
.git/  
├── hooks/  
│   ├── commit-msg  
│   ├── post-update  
│   ├── post-receive  
│   ├── pre-commit  
│   ├── pre-receive  
│   └── [...]
```

.git Directory Overview

```
.git/  
├── info/  
│   ├── exclude  
│   ├── attributes  
│   ├── refs  
│   └── sparse-checkout
```


.git Directory Overview

```
.git/  
├── objects/  
│   ├── 4d/  
│   │   └── 5fcadc293a348e88f777dc0920f11e7d71441c  
│   ├── 54/  
│   │   └── 59f2242e19a28f7fde64fcda2ba0f69cd338a8  
│   └── [...]  
├── info/  
│   └── packs  
└── pack/  
    ├── pack-8d329fd7e6a324224b53c99b3bdf8bf67a50f8b3.idx  
    └── pack-8d329fd7e6a324224b53c99b3bdf8bf67a50f8b3.pack
```

.git Directory Overview

```
.git/
├── refs/
│   ├── heads/
│   │   ├── master
│   │   └── bugfix/
│   │       └── 42-crash
│   ├── notes/
│   │   └── commits
│   └── remotes/
│       ├── origin/
│       └── master
├── stash
├── tags/
│   └── v1.0
```

.git Directory Overview

```
.git/
├── logs/
│   ├── HEAD
│   └── refs/
│       ├── heads/
│       │   ├── master
│       │   └── bugfix/
│       │       └── 42-crash
│       ├── remotes/
│       │   └── origin/
│       │       └── master
│       └── stash
```

References

- > SHA-1 ID (d25ecde05f2d52a5dc2dd022bf194a5a0522e127)
- > Short ID (d25ecde)
- > Branch Name (master)
- > Remote Branch (origin/master)
- > Tag Name (v1.0)
- > Path name (refs/heads/master)
- > HEAD
- > Relative to reference
 - > master^
 - > master@{2}
 - > master@{one.week.ago}

HEAD

> ref: refs/heads/master

or

> 0fcaf4244f90d93cbc7be8e670aeaa30288d24ae3

> All files in `.git/refs` contain a SHA-1 object ID.

Index

- > Binary file
- > Human readable: `git ls-files -stage`

```
100644 5bda33f260151f7b489cf747f9f2f186d77c9121 0    .gitignore
100644 a38a9a5d5a65d1aabc887f6fedf97ecff71af82c 0    LICENSE
100644 2902bfad3ea68bfccb60557fd7de7890e6eeb7ed 0    README.md
160000 0fcf4244f90d93cbc7be8e670aeaa30288d24ae3 0    wiki

100644 e69de29bb2d1d6434b8b29ae775ad8c2e48c5391 1    foo
100644 5716ca5987cbf97d6bb54920bea6adde242d87e6 2    foo
100644 257cc5642cb1a054f08cc83f2d943e56fd3ebe99 3    foo
```

Reference Logs

```
000000...000000 fecaf5...7c513a s3lph <s3lph@s3lph.me>↔  
1536434924 +0200      branch: Created from HEAD
```

```
fecaf5...7c513a b0390e...a6f537 s3lph <s3lph@s3lph.me>↔  
1536455366 +0200      commit: Fixed stuff.
```

```
b0390e...a6f537 2161ff...aeb311 s3lph <s3lph@s3lph.me>↔  
1536455387 +0200      pull: Fast-forward
```

```
2161ff...aeb311 37498f...8b86d1 s3lph <s3lph@s3lph.me>↔  
1536455840 +0200      merge: origin/dev: fast-forward
```

Objects

```
TYPE      = "blob" / "tree" / "commit" / "tag"  
LENGTH   = 1*DIGIT  
VALUE    = *OCTET  
OBJECT   = TYPE 0x20 LENGTH 0x00 VALUE
```


git cat-file



tpholland, "I'm just looking for something.", CC BY 2.0.

- > `$ cat .git/objects/0f/cf424... | inflate`
 - > Cumbersome
- > `$ git cat-file -p ref`
 - > Value only
 - > Pretty-printed
 - > Easier to use

Tree

```
100644 blob 5bda33f260151f7b489cf747f9f2f186d77c9121 .gitignore
100644 blob a38a9a5d5a65d1aabc887f6fedf97ecff71af82c LICENSE
100644 blob 2902bfad3ea68bfccb60557fd7de7890e6eeb7ed README.md
040000 tree 6229aa2ed3e0885e59562714bacf8b0247f851ee src
160000 commit 0fcf4244f90d93cbc7be8e670aeaa30288d24ae3 wiki
```

> IDs and types stored in binary form

Commit

```
tree bcbab69addb35bd6563e280afd5a2afa26c59633
parent a620e219a45cb93ca6bdc95209f17c4e092d0ad3
parent 9b1d8eb61c329fdbeb6cf8b043a80fd7eed9e13d
author s3lph <s3lph@s3lph.me> 1549599264 +0200
committer s3lph <s3lph@s3lph.me> 1549599264 +0200
gpgsig -----BEGIN PGP SIGNATURE-----
-----END PGP SIGNATURE-----
```

Merge branch 'bugfix/42-crash' into 'master'

Fix #42: Crash.

Tag

```
object a620e219a45cb93ca6bdc95209f17c4e092d0ad3
type commit
tag v1.0
tagger s3lph <s3lph@s3lph.me> 1549599264 +0200
```

Release 1.0

Notes

```
tree 9f8ba04eb8de0d68a644e38ccba25b8607c625f1
parent 3392844630ae2a682f13dca350504b915763ea87
author s3lph <s3lph@s3lph.me> 1560454563 +0200
committer s3lph <s3lph@s3lph.me> 1560454563 +0200
```

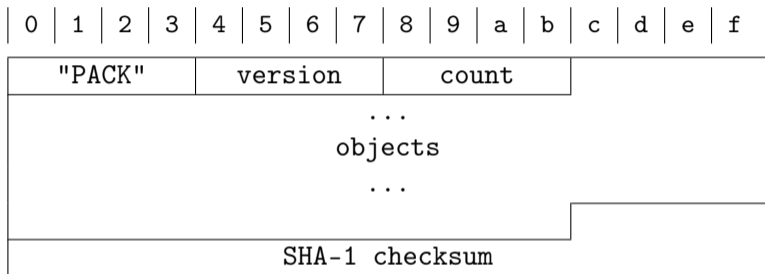
Notes added by 'git notes add'

```
100644 blob 7f5d90...d8ecc6 8ba9be5dc322440755291b298d4c12b2ad36104a
100644 blob d90bda...1813c6 d25ecde05f2d52a5dc2dd022bf194a5a0522e127
```

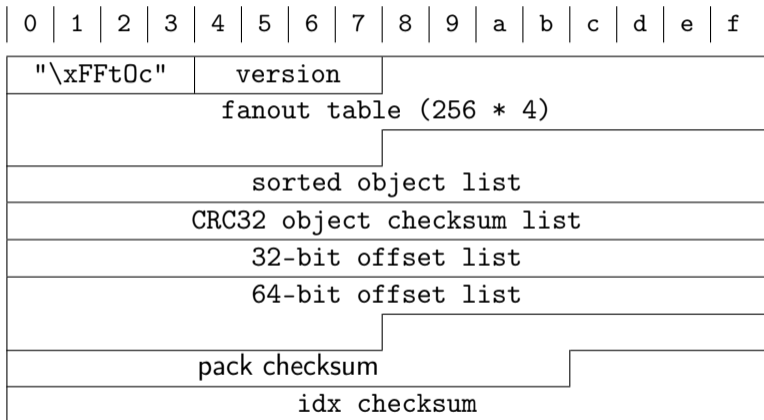
Packfiles

- > Lots of small files
- > Lots of similar files (revisions)
- > Packfiles
 - > Objects bundled together
 - > Similar objects stored as delta
 - > 2 new types: `OFS_DELTA` and `REF_DELTA`
- > Pack Index
 - > Detached lookup table
 - > Search tree

Packfile Layout



Packfile Index Layout



Garbage Collection

- > Prune old abandoned objects
- > Prune reflog
- > Optimize packfiles
- > Run automatically on `git push`
- > Manually: `git gc`

Push

- > `git send-pack`
- > SSH → `git-receive-pack`
- > `$ ssh -x gitlab.com 'git-receive-pack me/myrepo.git'`

Push

```
< 009d6a3c56...bf939c refs/heads/master%00report-status delete-refs ...  
< 003e99ee27...72e1f6 refs/heads/dev  
< 0000
```

Push

```
< 009d6a3c56...bf939c refs/heads/master%00report-status delete-refs ...
< 003e99ee27...72e1f6 refs/heads/dev
< 0000

> 006e99ee27...72e1f6 86e0b3...3b11b8 refs/heads/dev%00report-status
> 006c000000...000000 99ee27...72e1f6 refs/heads/bugfix/42-crash
> 0000
```

Push

```
< 009d6a3c56...bf939c refs/heads/master%00report-status delete-refs ...
< 003e99ee27...72e1f6 refs/heads/dev
< 0000

> 006e99ee27...72e1f6 86e0b3...3b11b8 refs/heads/dev%00report-status
> 006c000000...000000 99ee27...72e1f6 refs/heads/bugfix/42-crash
> 0000

> (packfile)
```

Push

```
< 009d6a3c56...bf939c refs/heads/master%00report-status delete-refs ...
< 003e99ee27...72e1f6 refs/heads/dev
< 0000

> 006e99ee27...72e1f6 86e0b3...3b11b8 refs/heads/dev%00report-status
> 006c000000...000000 99ee27...72e1f6 refs/heads/bugfix/42-crash
> 0000

> (packfile)

< 000eunpack ok
```

Push via HTTP

```
> GET /me/myrepo.git/info/refs?service=git-receive-pack HTTP/1.1

< 009d6a3c56...bf939c refs/heads/master%00report-status delete-refs ...
< 003e99ee27...72e1f6 refs/heads/dev
< 0000

> POST /me/myrepo.git/git-receive-pack HTTP/1.1

> 006e99ee27...72e1f6 86e0b3...3b11b8 refs/heads/dev%00report-status
> 006c000000...000000 99ee27...72e1f6 refs/heads/bugfix/42-crash
> 0000
> (packfile)
```

Fetch

- > `git fetch-pack`
- > SSH → `git-upload-pack`
- > `$ ssh -x gitlab.com 'git-upload-pack me/myrepo.git'`

Fetch

```
< 00dfca82a6...763949 HEAD%00symref=HEAD:refs/heads/master ...
< ...
< 0000

> 003cwant ca82a6...763949%00ofs-delta ...
> 0032have 085bb3...06e7e7
> 0009done
> 0000

< (packfile)
```

Fetch via HTTP

> GET /me/myrepo.git/info/refs?service=git-upload-pack HTTP/1.1

< 00dfca82a6...763949 HEAD%00symref=HEAD:refs/heads/master ...

< ...

< 0000

> POST /me/myrepo.git/git-upload-pack HTTP/1.1

> 003cwant ca82a6...763949%00ofs-delta ...

> 0032have 085bb3...06e7e7

> 0009done

> 0000

< (packfile)

Plumbing Commands Overview

- > `git cat-file` → Pretty-print a Git object
- > `git hash-object` → Create a blob object
- > `git mktree` → Create a tree object
- > `git commit-tree` → Create a commit object
- > `git mktag` → Create a tag object
- > `git update-index` → Modify (add, delete, ...) index
- > `git pack-objects` → create packfile
- > `git unpack-objects` → unpack packfile
- > ...



Questions?

s3lph@s3lph.me

@s3lph@chaos.social