

G O T

## Game of Trees Daemon

Stefan Sperling <stsp@openbsd.org>

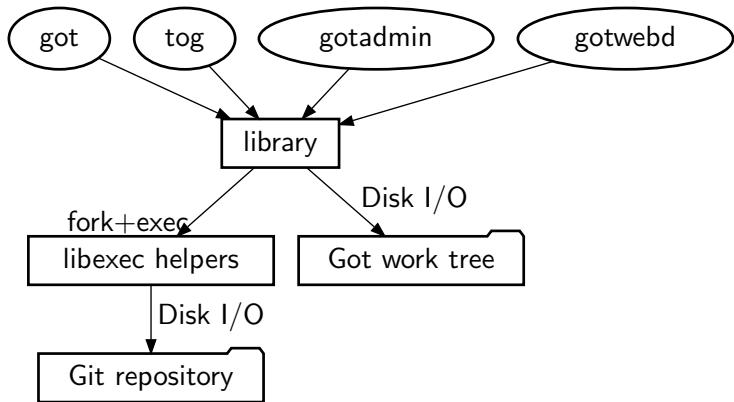
FOSDEM 2023

# What is Game of Trees?

Game of Trees (Got) is a distributed version control system.

- in development since November 2017
- compatible with Git repositories and Git network protocol
- uses OpenBSD pledge and unveil
- ISC licence
- -portable version for Linux, MacOS, {Free,Net,Dragonfly}BSD

# Game of Trees Clients Design



# got – command line interface

## Commands:

- import, clone, fetch, checkout, update
- log, diff, blame, tree, status
- ref, branch, tag
- add, remove, patch, revert, commit, send
- cherrypick, backout, rebase, histedit
- integrate, merge, stage, unstage, cat

Online documentation: <https://gameoftrees.org/got.1.html>

# gotadmin – repository administration

Commands:

- `init`, `info`, `pack`, `indexpack`, `listpack`, `cleanup`

Online documentation:

<https://gameoftrees.org/gotadmin.1.html>

# tog – ncurses-based interactive repository browser

```
commit e982275c90939242a47626658c419c91fdffea50 [47/6
2022-12-30 jsg add history for getpid(2) and ge
2022-12-30 cheloha pause.3: miscellaneous rewrites,
2022-12-30 cheloha accton.c: add missing $OpenBSD$
2022-12-30 kettenis Do not send (normal) packets bef
2022-12-30 kettenis Add chip name for new revision o
2022-12-30 jca Actually hide the clang-15 worka
2022-12-30 jca Neuter zlib fatal warnings when
2022-12-30 kevlo Support FDI FT232R. The upper
2022-12-29 cheloha sparc64: pull retry logic out of
2022-12-29 schwarze HISTORY: clarify that unmount(2)
2022-12-29 kettenis Avoid doing cache flush/invalida
2022-12-29 kettenis Use an iorw fence since we're sy
2022-12-29 jmc capital letter at sentence start
2022-12-29 guenther FORK_SIGHAND and the handling of
2022-12-29 jsg profil(2) first appeared in fift
2022-12-29 jsg fcntl(2) first appeared in Syste
2022-12-29 guenther Add ktrace struct tracepoints fo
2022-12-28 jmc {en,de}queing -> {en,de}queuing;
2022-12-28 jmc spelling fixes; from paul taglia
2022-12-28 cheloha iostat(8): implement periodic di
2022-12-28 cheloha systat(1): vmstat: dinfo(): comp
2022-12-28 cheloha microtime.9: rewrite description
2022-12-28 tb style(9) for includes
2022-12-28 jsg the S in CSRC is Science not Sci
2022-12-28 kettenis The RISC-V architecture support
2022-12-28 tb Properly ignore comments in geof
2022-12-28 kettenis Handle clock that can't be gated
2022-12-28 asou Change space character to TAB.
2022-12-28 kn Make wait_reorder_libs() honour
2022-12-28 jsg success -> success
2022-12-28 jca Ansify, fixes clang 15 -Wdepreca
2022-12-28 yasuoka Fix the word in the copyright.
2022-12-27 jmc spelling fixes; from paul taglia
2022-12-27 kettenis Newer device trees for boards ba

[1/32] diff 879f4f875f88f30442229e1c78aafdb38dce13f8 e982275c90939242a47626658c4
commit e982275c90939242a47626658c419c91fdffea50
from: jmc <jmc@openbsd.org>
date: Thu Dec 29 06:49:34 2022 UTC

capital letter at sentence start;
M share/man/man9/Fork1.9 | 2+ 2-


1 file changed, 2 insertions(+), 2 deletions(-)

commit - 879f4f875f88f30442229e1c78aafdb38dce13f8
commit + e982275c90939242a47626658c419c91fdffea50
blob - 4a83c1564ff09f4f8a2c85acb980106acd3e2c27
blob + e592b902dfe57fcb578e49999da87214299ba22e
--- share/man/man9/fork1.9
+++ share/man/man9/fork1.9
@@ -1,4 +1,4 @@
-.\ " $OpenBSD: fork1.9,v 1.31 2022/12/29 06:10:54 guenther Exp $
+.\ " $OpenBSD: fork1.9,v 1.32 2022/12/29 06:149:34 jmc Exp $
.\ " $NetBSD: fork1.9,v 1.3 1999/03/16 00:40:47 garbled Exp $
.\ "
.\ " Copyright (c) 1998 The NetBSD Foundation, Inc.
@@ -106,7 +106,7 @@ the new thread will begin execution by calling
ptrace(PT_TRACE_ME, 0, 0, 0) had been invoked in the child.
.El
.Pp
-the new thread will begin execution by calling
+the new thread will begin execution by calling
.Fa func ,
which must not be
.Dv NULL .

(END)
```

Supports log, diff, blame, tree, and ref views

# gotwebd – repository viewer for web browsers



The logo for Game of Trees (GOT) features the letters 'G', 'O', and 'T' in a hand-drawn, sketchy font. The 'O' is a circle with a smiling face, including two dots for eyes and a curved line for a mouth.

Game of Trees Projects		
Project	Description	Last Change
diff.git	New diff implementation	2 months ago
<a href="#">summary</a>   <a href="#">briefs</a>   <a href="#">commits</a>   <a href="#">tags</a>   <a href="#">tree</a>   <a href="#">rss</a>		
got-portable.git	Portable version of Game of Trees (Linux, ...)	20 hours ago
<a href="#">summary</a>   <a href="#">briefs</a>   <a href="#">commits</a>   <a href="#">tags</a>   <a href="#">tree</a>   <a href="#">rss</a>		
got-www.git	The gameoftrees.org web site	20 hours ago
<a href="#">summary</a>   <a href="#">briefs</a>   <a href="#">commits</a>   <a href="#">tags</a>   <a href="#">tree</a>   <a href="#">rss</a>		
got.git	Game of Trees	21 hours ago
<a href="#">summary</a>   <a href="#">briefs</a>   <a href="#">commits</a>   <a href="#">tags</a>   <a href="#">tree</a>   <a href="#">rss</a>		

Shows commits, diffs, trees, blobs, and tags (with RSS feed)

# gotd – Game of Trees Daemon

## Git repository server implementation

- in development since September 2022
- can be installed on OpenBSD -current: `pkg_add gotd`
- not yet available in -portable



# gotd – Game of Trees Daemon

Use cases:

1. host Git repositories for open source projects or private use
  - each project should host their own server instance
2. host public read-only Git repository mirrors
  - let anonymous users fetch source code over SSH,<sup>1</sup> authenticating the server but not the client

---

<sup>1</sup>like AnonCVS: <https://www.openbsd.org/papers/anoncvsv-paper.pdf>

## gotd live on the internet

git clone URL for OpenBSD version of Got:

```
ssh://anonymous@got.gameoftrees.org/got.git
```

git clone URL for -portable version of Got:

```
ssh://anonymous@got.gameoftrees.org/got-portable.git
```

Host key fingerprints:

- SHA256:aDX3rrQNDfIA5lyXIbynY+goiomgR4Cdx5j4qmWq26I (DSA)
- SHA256:kmeRVbw2BLak1iZJIv6/AkNhGqW97WIu58SeH8kUf1s (ECDSA)
- SHA256:jvQMqNSKE+p7roYA3Tl8/giZyOUqL7emXS+1JT7KMq4 (ED25519)
- SHA256:q/CWVAukwojxNL3srvWhfHZX00t+e01IhmnwFAEhE6o (RSA)

Same info also on <https://gameoftrees.org/code.html>

# Git protocol excursion

Git client connection setup during `git clone`:

- Git client logs in via SSH, and runs:
- `$SHELL -c 'git-upload-pack /git/got.git'`
- `git-upload-pack` speaks Git protocol on standard output

Let's look at an example Git protocol trace, as shown by "`got clone -v`", to see what happens next.

## Git protocol excursion

The server sends the first protocol message, listing one of the available Git branches and supported Git protocol capabilities:

```
readpkt: 122:  c4d35c5bb4f936d0f96bb62d234001b68dc33089 HEAD[0x00] \  
  agent=got/0.84-current ofs-delta side-band-64k \  
  symref=HEAD:refs/heads/main[0x0a]
```

Each Git protocol message is a “packet-line”, a simple length + data framing scheme (above message length is 122 bytes).

## Git protocol excursion

The server sends more messages, listing all available Git branches

```
readpkt: 57:    c4d35c5bb4f936d0f96bb62d234001b68dc33089 \  
               refs/heads/main[0x0a]
```

The list is terminated by a “flush-packet” (length zero):

```
readpkt: 0:
```

# Git protocol excursion

The client sends a similar list of branches it wants:

```
writepkt: 0061: want c4d35c5bb4f936d0f96bb62d234001b68dc33089 \  
    agent=got/0.84-current ofs-delta side-band-64k[0x0a]
```

This list is also terminated by a flush packet:

```
writepkt: 0000
```

# Git protocol excursion

The client could now send “have” lines listing branch tips it has stored locally.

But during a fresh clone, the client immediately sends its final message:

```
writepkt: 0009: done[0x0a]
```

## Git protocol excursion

The server sends a NAK, indicating no common ancestor commits were found, as expected during a fresh clone.

```
readpkt: 4:      NAK[0x0a]
```

Common ancestors found would be indicated by “ACK” messages. They would allow for reducing the size of the pack file sent next, and stop the client sending “have” lines.



## Git protocol excursion

The server creates a pack file containing objects reachable via the requested commits, and sends progress output:

```
server: 5062 commits colored, 28448 objects found, deltify 100%
```

# Git protocol excursion

The client receives the pack file and creates an index for it:

```
8.9M fetched; indexing 100%; resolving deltas 100%  
Fetched 7a5fa8a611c8c665a052f637510bdceaaf7bc38f.pack
```

The client creates references which point at fetched branch tip commits, and the repository is ready for use.

# Git protocol excursion

The “git push” case:

- client receives initial reference list
- client proposes reference updates
- client sends a pack file
- server indexes and verifies pack file
- server accepts or denies reference updates

Details:

`Documentation/gitprotocol-pack.txt` in Git's source tree

# OpenBSD-style multi-process program

Common design patterns in OpenBSD daemons:

- ensure system-call domain-separation between processes
  - e.g. either network, or file access, or forking; Do not mix
  - use `pledge(2)` to help guide and enforce this
- `unveil(2)` restricts which files can be seen, read, or written
- `fork+exec` to create processes with unique memory layout
  - usually starting the same executable again with different command line flags indicating desired child process behaviour
- inter-process communication messages over pipes
- pass file-descriptors to make files and network sockets accessible to processes which cannot open them

# gotd – Game of Trees Daemon

Implements a Git server as a multi-process program.

- network connections via SSH only
- Git user accounts are Unix system shell accounts
- special-purpose shell can prevent access to arbitrary commands<sup>2</sup>
- access permissions are set per repository

---

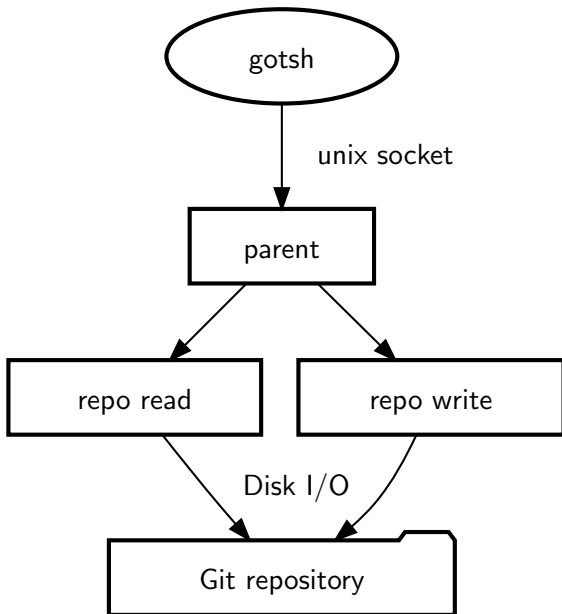
<sup>2</sup>similar to `git-shell` but more restricted and with less features

## /etc/gotd.conf – gotd configuration file

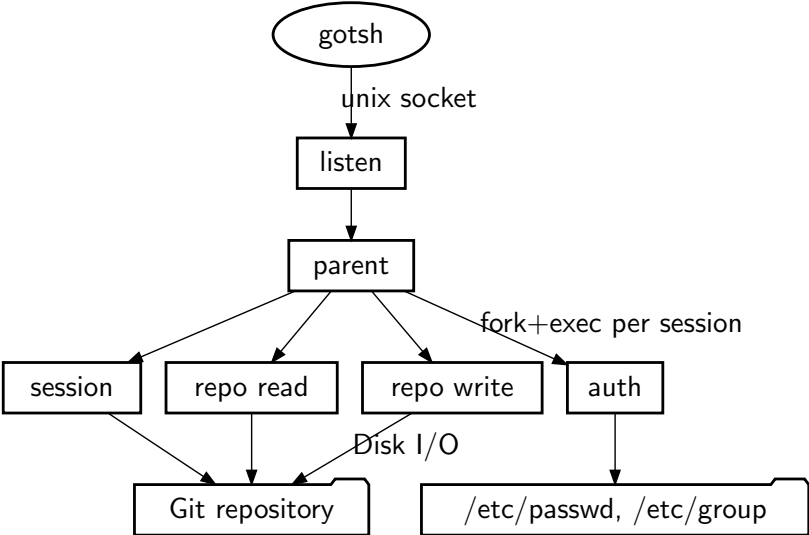
gotd requires a configuration file in order to run:

```
repository "test" {  
    path "/var/git/test.git"  
    permit rw :developers  
    permit ro anonymous  
}
```

# First Working Implementation



# Current Implementation






## gotd – parent process

- drop root privileges to user `_gotd`
- `pledge("stdio proc exec sendfd recvfd")`
- `unveil(argv[0], "x");`

parent


## gotd – listen process

- open unix socket `/var/run/gotd.sock`
- drop root privileges to user `_gotd`
- `pledge("stdio sendfd unix")`
- `unveil("/", "");`
- enforce per-UID connection limit




# gotsh – Game of Trees Shell

- `pledge("stdio recvfd unix")`
- connect to unix socket
- `pledge("stdio recvfd")`
- translate between packet-lines and internal messaging




## gotd – authorization process

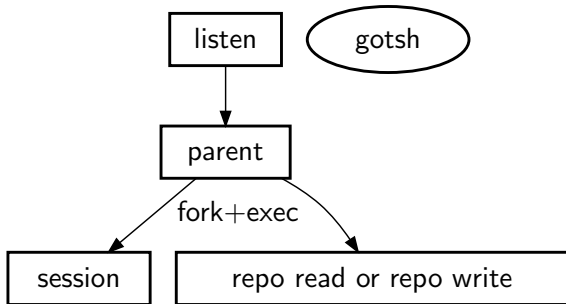


## gotd – authorization process

- `pledge("stdio getpw recvfd unix")`
- `unveil("/", "")`
- match repository access rules against user and groups
- report authorization result to parent and exit




# gotd – starting session and repo processes



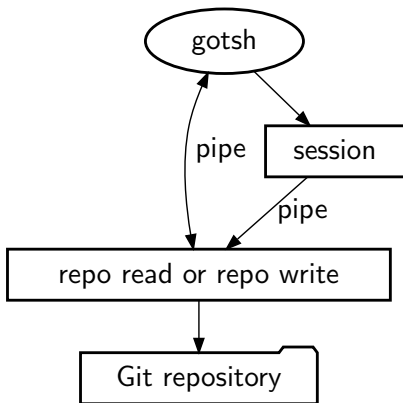
## gotd – session process

- `pledge("stdio rpath wpath cpath sendfd fattr flock")`
- `unveil(repo_path, "rwc"); unveil("/tmp", "rwc")`
- Git protocol state machine, driven by gotsh
- create temporary files needed by repo process
- install pack files uploaded by clients and update refs



## gotd – repo read or repo write process

- `pledge("stdio rpath recvfd")`
- `unveil(repo_path, "r")`
- repo read: create pack and stream it to gotsh
- repo write: receive pack from gotsh and create pack index





## Pending implementation improvements

- verify content of uploaded pack files
- parse configuration file just once on startup, not whenever a new child process starts
- split “session” process to avoid write access when not needed
- revisit Git protocol state machines in gotsh and “session”:

```
gotd[74429]: received flush-pkt from uid 1002  
last message repeated 386788 times
```

(above state machine bug has been fixed but there should be more...)

## Planned features: “pre-commit” checks

Built-in checks which can be enabled in the configuration file:

- enforce a configurable blob size limit
- deny branch history rewriting
- deny creation/deletion of selected references
- hide selected references from clients
- deny addition of binary files
- deny addition of merge commits

## Planned features: commit notifications

- commit email notification, plaintext SMTP to localhost
- send general-purpose HTTP request as commit notification
  - with format-string expanding commit info in URL/body
  - can be used to trigger arbitrary post-commit hook scripts

## Other planned features

- keep track of available repository disk space and fail gracefully
- teach “gotadmin cleanup” to remove redundant pack files
- add support for fast-import/fast-export to “gotadmin”
- implement Git’s SHA256 object ID support in repository and network protocol, and enable it by default
- server-side rebasing<sup>3</sup> to keep linear history without forcing clients to do trivial rebasing before sending commits

---

<sup>3</sup>Not the same as a similarly named planned feature for Git.

Thank you for listening! Got questions?



<https://gameoftrees.org>