

SWARMFILE

Documentation

Version 0.3.10

Generated 2026-10-09 from <https://swarmfile.com/docs>

The complete consumer documentation: guides, admin and IT, reference, and the command-line tools. The always-current version lives at <https://swarmfile.com/docs>. This PDF was generated from the same markdown the site renders.

Contents

Guides

- Getting Started
- Coming from Git
- Moving from Perforce
- Migrating Existing Data
- Deployment Topologies
- Working with Files
- Working Offline (Pack & Go)
- Multiple Mounts
- File Previews
- Search
- Notifications & Inbox
- Sync Exclusions
- Performance & Disk Usage
- Version Control
- Branches & Merging
- Trash, History & Rollback
- Verifiable History
- Git and Swarmfile
- Cloning a Project with Git
- Git-LFS
- Coding Agents
- Runner (Headless CI)
- Hosted CI
- Automating Swarmfile
- Sharing & Collaboration
- Worksharing (no plugin)
- RFIs & Submittals
- Public Projects & Org Profiles
- Publishing Releases
- Troubleshooting
- Uninstall & Clean Reinstall

Admin & IT

- Organizations, Projects & Members
- Identity
- LDAP Directory Sync
- Permissions
- Self-Hosted Seed Nodes
- Security
- Security Architecture (Whitepaper)
- Trust & Compliance
- End-to-End Encryption Setup
- Dedicated Storage Isolation
- Bring Your Own Storage
- Data Residency
- Billing & Plans
- Operations
- Webhooks
- Deploying to Your Team
- Data Portability & Offboarding
- Branch Mirror to S3

Reference

- Filesystem Compatibility & Conformance
- Supported Platforms & System Requirements
- Release Channels & Updates
- Diagnostics & Error Reference
- Environment Variable Index

Engine Config File (config.json)

Network Requirements

Storage Format

Telemetry & Data Collection

Known Limitations

CLI Reference

swarmfile

swarmfile-doctor

swarmfile-migrate

swarmfile-search

swarmfile-seed

swarmfile-brlock

swarmfile-runner

swarmfile-verify-history

swarmfile-lfs-transfer

git-remote-swarmfile

Getting Started

Swarmfile mounts your team's project storage as a real local drive. It is not a sync folder — nothing downloads to your machine until you actually open a file. Browsing a folder with tens of thousands of files costs zero disk space. Opening a 200GB file streams only the bytes your app reads.

Installing the client

Download the installer for your platform from [Downloads](#).

- **macOS** — a signed and notarized universal `.dmg` / `.pkg` that runs on Apple Silicon and Intel. The client uses FUSE-T, so there is no kernel extension: no Recovery-mode "Reduced Security" step, no rebooting to approve a kext, no security prompts beyond the normal Gatekeeper install.
- **Windows** — a Setup bundle that chains the required dependencies (WinFsp, the VC++ Redistributable, and WebView2) automatically, so you don't need to install them separately. Both ARM64 and x64 builds are supported.
- **Linux** — a `.deb` package (Debian/Ubuntu). Need an `.rpm` (Fedora/RHEL/openSUSE)? [Get in touch](#) and we'll provide one.

The `.pkg`, MSI, and `.deb` installers also put the `swarmfile` command line on your `PATH`, along with the standalone tools (`swarmfile-lfs`, `swarmfile-doctor`, `swarmfile-migrate`, `swarmfile-search`, `swarmfile-verify-history`, `swarmfile-seed`, `swarmfile-runner`) and the `git-remote-swarmfile` helper that backs `git clone swarmfile://` — see [CLI: swarmfile](#). There is no separate CLI download. On macOS, the `.pkg` is what installs those `PATH` entries: the `.dmg` installs the Desktop App, whose bundle carries the same tools but doesn't link them into `/usr/local/bin`, so a `.dmg`-only install needs one `.pkg` (or [Diagnostics → Reinstall](#)) run to put the CLI on `PATH`. (`swarmfile-brlock` is a developer reference client built from source, not part of the desktop install.)

Windows: the SmartScreen prompt

The first time you run the downloaded installer, Windows may show a blue "**Windows protected your PC**" box from Microsoft Defender SmartScreen. This is expected: the `Setup.exe` and `.msi` installers are Authenticode-signed, but SmartScreen warns until the signing identity has built up download reputation, and the `.zip`'s tools plus the standalone `swarmfile-doctor` download are not code-signed at all. Either way, it does not mean the download failed or was tampered with.

1. Click **More info** — the small link, not **Don't run**; the **Run anyway** button appears underneath.
2. Click **Run anyway**.
3. If Windows asks for permission, choose **Yes**, then finish the installer as usual.

You see the prompt when you run a freshly downloaded installer; automatic updates inside the app don't show it. A scripted mass deployment (`msiexec`, GPO/SCCM/Intune) never sees it — see [Deploying to Your Team](#).

Signing in

Open the client and sign in with your email and password, or with your organization's SSO if your IT team has configured an identity provider. The default form is email/password against the built-in identity provider; for SSO, click "Sign in with SSO" and enter your **organization's slug** (the short name in your org's Swarmfile URL) — the client looks it up and shows a "Continue with <your IdP>" button if that org has SSO configured. There's no email-domain auto-detection: you need to know your org's slug (or use the direct per-org sign-in link your admin can share) rather than just entering your email and having it figured out for you. An organization on Starter or above can also require an authenticator app, a verified email, or both before it admits members; if yours does, you're prompted to enroll before enforcement starts. If your account has an authenticator app, sign-in asks for its current six-digit code after your password (enter a recovery code instead if you don't have the app), and a wrong code can simply be retried. See [Identity](#).

Switching to a **different account** while the app is running is refused (`account_switch_requires_restart`): the local cache identity is reconciled only at startup, so that one user's unsent work can never be uploaded as another's. Sign in with the other account and restart when the Desktop App offers it — the requested sign-in is kept, and the restart completes the switch; any unsent work from the previous account is quarantined on your machine rather than lost.

New to Swarmfile? Starter and Pro start with a **7-day free trial** covering up to 2 seats (signing up past 2 seats starts paid), and a no-card Free plan covers public projects up to 1 GiB — see [pricing](#).

Creating a project

If your organization doesn't have the project yet, create one from the Desktop App or the web app. You can start from a **template** — Blank, Media & post, Software (git repository), Revit worksharing, or Geospatial — which pre-sets sensible defaults (encryption tier, commit mode, default-branch protection). Media & post, Revit worksharing, and Geospatial also seed a starter folder layout plus an ignore file so everyone begins from the same shape; Software (git repository) seeds just the ignore file, and Blank starts empty. A template that seeds an ignore file can't be used on an **end-to-end encrypted** project — the hub holds no key, so the create is refused; start Blank there, or create the project on a plaintext or managed tier and keep the seed.

Two creation choices are fixed and worth a moment's thought. The **initial branch** is the name of the project's default branch: most projects keep `main`; choose another (for example `develop`) only if your tooling expects it. It can't be changed later, it's the branch a fresh clone checks out, and the web app opens the project on it — so `git clone` of a project whose default is `develop` lands on `develop`, and its Files, Commits and Runs views start there too. The **encryption tier** — managed (the default), end-to-end, or plaintext on a public project — is also fixed at creation; upgrading later doesn't retroactively encrypt an existing project. See [Security Architecture](#).

To inspect a project's effective settings — every option's value and where it came from — run `swarmfile project show`; change them with `swarmfile project set <key=value>` (see the [CLI reference](#)).

Mounting a project

Once signed in, the Desktop App shows the projects you have access to. Pick one and it appears as a drive on your machine — a drive letter on Windows, a mount point on macOS and Linux. The exact location is in the Desktop App's header: click the **Status pill** to open the panel showing the mount point and sync state. It's assigned at mount time and can vary between machines (for example, if a default drive letter is already taken, Swarmfile picks the next free one).

From here, the mount behaves like any other local filesystem. Open it in Finder, Explorer, or your file manager, or point a native app directly at a file inside it.

First open

The first time you open a file, only the bytes your application actually reads are streamed down. Scrubbing 30 seconds into a 200GB video timeline streams roughly 30 seconds of footage, not the whole file. Browsing folders, previewing thumbnails, and listing directories don't trigger downloads at all — that metadata is served directly.

Presenting the Desktop App

Cmd/Ctrl+Shift+P toggles **presentation mode**, the on-camera switch for demos and screen recordings. Text scales to the largest size so a 1080p capture stays readable after compression, and names, emails and organization names are replaced with stable pseudonyms ("Person AB") so a customer's details can't slip into a video. The switch is local to your machine — it never changes what teammates see.

Prefer larger text all the time, no recording involved? The Desktop App's **Settings → Text size** offers Compact, Default and Large steps for that computer (Cmd/Ctrl + and – step it, Cmd/Ctrl 0 returns to Default); presentation mode always uses the largest step.

For multi-machine demos, presentation mode can also show a small always-on-top clock overlay, with the time synced from the network so machines shown side by side agree.

Where to go next

- [Working with Files](#) covers browsing, saving, presence, and locking in more depth.
- New to Swarmfile's own version control? [Version Control](#) covers changesets, branches and rollback, and [Branches & Merging](#) covers working an option in isolation and landing it.
- Working in git? [Clone a Project with Git](#) covers `git clone swarmfile://` and `git push` against a project's git view, and [Git and Swarmfile](#) lays out the three ways the two fit together. Starting from a repository that already exists elsewhere? `swarmfile git import <url>` creates the project from it in one command — see the [import reference](#).
- If you're moving data in from an existing NAS or file server, see [Migrating Existing Data](#).
- Rolling this out to a whole team rather than installing by hand? See [Deploying to Your Team](#).

Coming from Git

If you know Git, you already know most of Swarmfile. The vocabulary is deliberately familiar — `commit`, `branch`, `merge`, `log`, `revert`, `tag` all mean what you expect. This page maps your Git habits onto Swarmfile and, more importantly, flags the few places where the mental model genuinely differs.

Swarmfile is its own tool, not a Git front-end. But the instincts carry over, and where a Git command has no equivalent, typing it anyway prints a one-line pointer instead of an error — so you can learn by doing.

The one big shift: mount instead of clone, and no save step

This is the whole difference in two sentences:

- **Your project is a live mounted drive, not a cloned copy.** For day-to-day work there's nothing to `clone`, `pull`, or `push`. Files stream from the cloud on demand; you open and save them like any local file, and other people's changes appear as you read. (When you genuinely want a git repository of a project — for git tooling, a build that expects one, or an offline archive — an owner can turn on git access and you can `git clone` it, and even `git push` commits back to it. The drive stays where day-to-day work happens.)
- **Saving is automatic.** There's no `add` / `commit` / `push` ceremony to make your work durable and shared — the moment you save, it starts syncing. Git fuses *saving* and *naming* into one mandatory step; Swarmfile splits them, so **naming a commit is optional, and you can do it whenever — even after the fact.**

Everything below follows from those two facts.

Cheat sheet

In Git	In Swarmfile	Notes
<code>git clone <url></code>	<i>(nothing)</i> — just mount · or <code>git clone swarmfile://<org>/<project></code>	Projects appear as a mounted drive once the app is running; <code>swarmfile status</code> shows where it's mounted. If the project has git access turned on, <code>git clone swarmfile://...</code> gives you a real git repository of it — see Cloning a Project with Git .
<code>git pull</code> / <code>git fetch</code>	<i>(nothing)</i> — or <code>git fetch</code> in a <code>swarmfile://</code> clone	The drive is live — a teammate's changes appear as you open files. <code>swarmfile log</code> shows recent commits. Note <code>swarmfile fetch</code> is not Git's fetch: it pulls a byte range of one file (including one still uploading).
<code>git push</code>	<i>(nothing)</i> — or <code>git push</code> from a <code>swarmfile://</code> clone	Saves stream to the cloud automatically. <code>swarmfile status</code> shows what's still uploading. A git clone of a git-enabled project can push commits back , fast-forward only.
<code>git add <path></code>	<code>swarmfile add</code> (opens a changelist if none)	No staging for everyday work — just save. <code>add</code> is idempotent: it opens a changelist so subsequent saves stage into it (and this mount stays in staged mode until you <code>commit</code>), and it exits 1 on a path that doesn't resolve. For a deliberate multi-file batch that lands all at once, see Version Control .
<code>git commit -m "..."</code>	<code>swarmfile commit -m "..."</code>	Names a specific save. You don't <i>have</i> to — unnamed autosave commits happen for you as you work.
<code>git commit --amend -m "..."</code>	<code>swarmfile commit --amend -m "..."</code>	Renames the most recent commit. Message-only and non-destructive — contents and history are untouched (Swarmfile never rewrites commits).
<code>git status</code>	<code>swarmfile status</code> (mount/engine) · <code>swarmfile changelist status</code> (staged) · <code>swarmfile diff</code> (branch vs base)	<code>swarmfile status</code> is <i>not</i> Git's working-tree status — it reports the mount, peers, and settings. There's no "modified but unsaved" state to show: saves sync immediately. Use <code>changelist status</code> for staged work and <code>conflicts</code> for blocked files.
<code>git log</code>	<code>swarmfile log</code>	Compact, <code>--oneline</code> -style feed.
<code>git show <commit></code>	<code>swarmfile show <N></code>	Whole-file: shows which files changed (A/M/D/R), not a line diff — Swarmfile versions binaries.
<code>git diff main...</code>	<code>swarmfile diff [<branch>]</code>	Previews what a branch would merge into its base — files added/modified/deleted/renamed, plus conflicts. Whole-file (no line diff).
<code>git diff A B</code>	<code>swarmfile diff <A> </code>	Files changed between two commits (each a number, <code>HEAD~N</code> , tag, or branch). Whole-file.
<code>git blame <file></code>	<code>swarmfile blame <path></code>	Who last changed the file + its version history with authors — one row per saved version, not per line.
<code>git switch <branch></code>	<code>swarmfile switch <branch></code>	Hot, in-process — no restart.
<code>git checkout <branch></code>	<code>swarmfile switch <branch></code>	<code>swarmfile checkout <branch></code> also works (it's a shim); <code>switch</code> is clearer.
<code>git checkout -b <name></code> / <code>git switch -c <name></code>	<code>swarmfile checkout -b <name></code> / <code>swarmfile switch -c <name></code>	Creates the branch (forked from your current one — or, like git, from <code>checkout -b <name> <start-point></code>) and switches to it in one step.
<code>git restore <path></code> / <code>git checkout <commit> -- <path></code>	<code>swarmfile restore <N> --path <path></code>	Rolls a file, folder, or the whole project back to a past commit, landed as one undoable commit.
<code>git revert <commit></code>	<code>swarmfile revert <N></code>	Lands a new commit undoing an old one. Non-destructive and itself undoable.
<code>git branch</code>	<code>swarmfile branch list</code>	
<code>git branch --merged</code>	<code>swarmfile branch merged</code>	Branches with nothing left to land (branch-vs-base diff empty); <code>swarmfile branch prune</code> archives them after confirming, and never archives a branch an open mount on this machine is working on.
<code>git add --resolved</code> (staging resolved conflicts)	<code>conflicts resolve <path> --auto, --winner S, --tool</code>	There's no index to stage into, and a conflicted file can't be saved until it's resolved — so a leftover-marker commit can't happen. Resolve, then saves sync. <code>--auto</code> already refuses (exit 2) if a merge leaves a real conflict.
<code>git tag <name></code>	<code>swarmfile tag create <name></code>	An immutable, named pointer to a commit — handy for "branch from here later."
<code>git merge</code>	<code>swarmfile merge</code> , or a merge request: <code>swarmfile mr create</code>	Whole-file merge with conflict detection. See Branches & Merging .
<code>git mergetool</code>	<code>swarmfile conflicts resolve <path> --tool</code>	Uses your git mergetool configuration, unless a <code>mergetool_cmd</code> is set in the engine's own config — that takes precedence over git config.
<i>(auto-merge)</i>	<code>swarmfile conflicts resolve <path> --auto</code>	Let the engine merge it (diff3 for text) with no external tool — exits 2 if a real conflict remains, so a script can fall back. <code>--all --auto</code> sweeps every conflicted file in one pass.
<code>git rm <path></code>	<i>(do it on the drive)</i>	Delete the file on the mounted drive; it syncs like any other save and lands in Trash.
<code>git mv <a> </code>	<i>(do it on the drive)</i>	Rename/move on the mounted drive. The entry keeps its id, so history, comments, and locks move with it.
<code>git worktree add</code>	<code>swarmfile mounts open <path></code>	One engine can hold several full mounts at once — the reason you reached for a worktree. Use <code>--project <id></code> for a different project (same org).
<code>git grep</code>	<code>git clone + git grep</code> , or <code>swarmfile-search</code>	Swarmfile's own search is metadata — names and paths, not file contents. To grep contents, work in a git clone of a git-enabled project; <code>swarmfile-search</code> (or the dashboard's search box) finds entries by name.
<code>git clean -fd</code>	<i>(nothing needed)</i>	Ignored paths are already kept off the hub. To remove content that synced before a rule existed, use <code>swarmfile ignore-clean</code> .

In Git	In Swarmfile	Notes
<code>git reflog</code>	<code>swarmfile log</code> / <code>swarmfile activity list</code>	History isn't rewritten, so there's no reflog — <code>log</code> lists commits, <code>activity</code> shows who changed what and when.
<code>git init</code> / <code>git remote</code>	<i>(nothing)</i>	There's no repo to initialize and no remotes — the drive is the project.
<code>.gitignore</code>	<code>.gitignore</code> / <code>.swarmfileignore</code>	Honored on sync. See Sync Exclusions .

`<N>` is a commit number — the `#N` shown by `swarmfile log`. Anywhere a commit is expected (`show`, `restore`, `revert`, `describe`, `checkout <commit>`) you can also use `HEAD`, `HEAD~N` / `HEAD^` (`N` commits back), a **tag name**, a **branch name** (its head), a changeset id, or a 64-hex **commit hash** — just like Git. A `seq: / head: / branch: / tag: / changeset: / hash:` prefix forces which kind is meant.

Log is per-branch, and shows a branch's *own* commits. `swarmfile log` (and `HEAD / HEAD~N`) reflect the branch this mount is on. Note one difference from Git: a branch's log shows the commits made *on that branch since it forked* — not the shared history from before the fork point. Switch to `main` to see the mainline history.

Your Git aliases work. If your `~/.gitconfig` defines `[alias] co = checkout`, `st = status`, etc., `swarmfile co` / `swarmfile st` expand the same way — Swarmfile reads your git aliases for any command name it doesn't already have.

Three habits to unlearn

1. **Don't reach for `add` / `commit` / `push` to "save and share."** Just save the file. It's durable and syncing before you'd have finished typing `git add`.
2. **Don't `clone` or `pull` to get to work.** The drive *is* the project, always current. `swarmfile status` tells you where it's mounted and what's in flight. A `git clone` of a git-enabled project is there for git tooling, archives and git-based workflows; it can push commits back, but it isn't the live drive.
3. **You don't have to name a commit when you make it.** Work all day; name the moments that matter afterward (see below). A blank autosave isn't a mistake — it's the default.

Naming work — now, or later

Because saving and naming are separate, you have three ways to put a name on history, and you can mix them freely:

- **Name as you go:** `swarmfile commit -m "Grade pass 2"` files your recent saves under that message.
- **Name after the fact:** `swarmfile describe <N> -m "v2 delivery to client"` names *any* past commit — including an autosave that landed with no message. `swarmfile commit --amend -m "..."` does the same for the most recent one. Both change only the message.
- **Mark a milestone:** `swarmfile tag create <name>` pins an immutable, memorable name to a commit.

In the desktop app and the web dashboard, the same thing is a pencil/**Add description** on any commit row — no command needed.

The two commit modes (30-second version)

- **Mode A (default, autosave):** saves land as commits automatically; a burst of edits becomes one commit, the next burst another. You do nothing. This is what an office worker who's never heard of Git gets, and it's fine.
- **Mode B (explicit):** open a *changelist*, stage a set of saves, and submit them as one atomic unit — the analog of the Git index, for work that only makes sense landing together.

They're not a setting you choose — open a changelist and you're in Mode B for those saves; otherwise you're in Mode A. Full detail in [Version Control](#).

What genuinely has no equivalent

- **git rebase / history rewriting** — Swarmfile history is an append-only server-side DAG of whole-file commits. It isn't rewritten. To put a name on the past, use `describe` or `tag`.
- **git reset (index-moving forms)** — `swarmfile reset --hard <rev>` does work, mapping to `restore <rev>`. But bare `reset`, `--soft` / `--mixed`, and `reset <path>` have no equivalent: saves sync immediately, so there's no unstaged state to move. To undo, land a `revert` (undo one commit) or a `restore` (roll a scope back to a point in time) — both move history *forward* rather than erasing it.
- **git stash** — actually maps onto Swarmfile's parked work: `swarmfile stash` runs `changelist park` (set staged work aside, kept on the hub), `stash pop` resumes it, `stash list / show` lists it, `stash drop` discards it. Parked work is a **set, not a stack** — `pop` restores *all* parked work for the branch, and a bare `drop` discards *all* of it, so it asks first; `stash list` numbers the entries and `stash drop stash@{1}` (or `drop 1`) discards just that one. In-progress saves are already auto-saved, so there's nothing lost either way.
- **File plumbing (`rm`, `mv`, `cp`, `mkdir`, `touch`) and `grep`** — these aren't CLI verbs. You do file operations on the mounted drive itself; for content search, work in a `git clone` of a git-enabled project. `swarmfile-search` finds entries by name, not file contents.

Type any of these (`swarmfile rebase`, `swarmfile rm`, `swarmfile worktree`, `swarmfile shortlog`, `swarmfile fsck`, ...) and Swarmfile will print the one-line reason and the thing to do instead. (`swarmfile add`, `swarmfile stash`, and `swarmfile reset --hard <rev>` don't just explain — they run the mapped command.)

The on-ramp: bring the git you already use (git-LFS)

For the side-by-side of every way git and Swarmfile meet — git-LFS, `git clone swarmfile://`, and the CLI's git-style commands — see [Git and Swarmfile](#).

You don't have to move a whole team at once. Step one is to keep your source in the git you use today — GitHub/GitLab/self-hosted — and route only the large binary assets to Swarmfile, where Swarmfile acts as a **git-LFS server**: the big files get real storage, dedup, encryption (on most upload paths — see the caveat below), and quota instead of bloating your git host.

1. In the dashboard, open your project → **git-LFS** tab → **Generate git-LFS credential**.

2. Commit the `.lfsconfig` it shows to your repo.
3. Run the setup commands it shows once in your working copy (they store the credential and force basic auth).

After that, `git lfs push / pull / checkout` move big files through Swarmfile with no change to your VCS or workflow. git-LFS works on **managed-key** and **unencrypted** projects; it isn't available on end-to-end encrypted projects (a plain `git-lfs` client can't hold the per-user key).

A few things worth knowing up front (all covered in the [Git-LFS](#) guide):

- **Pushes up to ~5 GiB** work with stock git-LFS; larger files — and chunked, resumable uploads — need the small transfer agent, which ships with the desktop app or installs standalone on a machine without one. Up to 256 GiB per object. Pulls normally need no agent; a very large object's first stock pull may be paced while the hub prepares it (the client retries), or refused on a deployment without the prepare queue — the mount reads it either way.
- **Locking** — `git lfs lock / unlock / locks` work, including for files that live *only* in git and were never written to the drive.
- **No lock-in** — `git lfs fetch --all` pulls every object back with stock git-LFS and no Swarmfile software, so you can repoint `.lfsconfig` at any other LFS host and leave. Your source was never in Swarmfile to begin with.
- **Onto the drive** — run `swarmfile-lfs reflect` after a push to have those pushed objects appear as first-class files on the mounted Swarmfile drive (reusing the uploaded blocks, no second copy).
- **Encryption posture** — git-LFS is the one surface where a managed project can hold plaintext by design: objects a stock `git-lfs` client pushes, and the desktop agent's very large uploads, are stored unencrypted. Anything that must be encrypted at rest belongs on the mounted drive. The [guide](#) has the full path-by-path table.
- **GitHub, a little more automatic** — on a GitHub remote, the git-LFS tab can install the Swarmfile GitHub App, open the `.lfsconfig` pull request for you, and reflect each push onto the drive without a local `reflect` (reflection still needs the transfer agent — a stock `git-lfs push` stores the object but doesn't place it on the drive).

The bridge is where you start, not where you stop. The fuller destination is making Swarmfile the home for the *whole* project: native version control — branch, merge, history and rollback — over a drive you mount and open files from, so nothing has to live behind a clone. `swarmfile git import <url>` takes an existing GitHub/GitLab/local repository there in one command — the source's default branch becomes the project's, commits keep their upstream ids, annotated tags convert to lightweight, and re-running it syncs new upstream commits into the same project (fidelity limits in [Known Limitations](#)). Move at your own pace; the LFS on-ramp buys you the storage win today.

Moving from Perforce

If your team runs Perforce (Helix Core) today, most of Swarmfile will feel familiar: changelists, exclusive checkouts of binary files, stream-style branches, review before a change lands. The big difference is that there's no workspace to sync. Your project is a **mounted drive**: files stream on demand, and a save is on its way to the team the moment you make it.

This page maps Perforce concepts onto Swarmfile, explains how to bring a depot across, and is direct about what doesn't carry over.

Concept map

Perforce	Swarmfile	Notes
Server / depot	Organization / project	A project is the unit of storage, history, and permissions.
Stream	Branch	Branches follow the streams model: a full working line over shared, content-addressed storage. Forking a branch copies no file data. See Branches & Merging .
Workspace (client spec) + <code>p4 sync</code>	The mounted drive — no sync	Files appear on demand and stay current as teammates save. There's no client view to maintain and no sync to run. <code>swarmfile switch <branch></code> moves a mount to another branch in place.
<code>p4 edit</code> / <code>p4 add</code> / <code>p4 delete</code>	Just open, save, create, or delete on the drive	Every change is tracked automatically.
Pending changelist	Changelist	Open one with <code>swarmfile changelist open -m "..."</code> ; saves stage into it until you submit. You can hold several at once and choose which is current, like <code>p4 ... -c</code> . See Version Control .
<code>p4 submit</code>	<code>swarmfile changelist submit</code> (or <code>swarmfile commit -m "..."</code>)	Applies atomically. Projects can also run in the default auto-landed mode, where each save lands as it's made and you name commits when it suits you.
Submitted changelist number	Commit number (#482)	<code>swarmfile log</code> , <code>swarmfile show 482</code> , <code>swarmfile diff</code> .
<code>p4 shelve</code> / <code>p4 unshelve</code>	Park / resume	<code>swarmfile changelist park</code> sets aside every open changelist on the mount (kept on the server); <code>changelist resume</code> brings it back. See Setting work aside .
Exclusive checkout (+l filetype, p4 lock)	Edit lock	<code>swarmfile lock <path></code> reserves a file for up to 30 days (7 by default); others can ask for it back with <code>swarmfile unlock-request</code> . See Entry locks .
Revit / SolidWorks exclusive open	Worksharing , no plugin	On Windows, the application's own exclusive open is enforced across machines. See Worksharing .
<code>p4 opened</code> / who has it	Presence	See who has a file open or locked, in the dashboard and the Desktop App.
Working disconnected	Pack & Go	Reserve a folder or project for offline editing and land the changes when you're back. See Working Offline .
<code>p4 undo</code> / rollback	<code>swarmfile revert</code> , <code>swarmfile restore</code> , <code>swarmfile rollback</code>	Revert a commit, restore a file, folder, or project to a past point — each as a new, undoable commit. See Trash, History & Rollback .
Labels	Tags	<code>swarmfile tag create v1.0</code> — an immutable name for a commit.
Protections table	Folder and file ACLs	Allow/deny read, write, or admin, inherited down the tree. See Permissions .
Helix Swarm review	Merge requests	Review on a branch before it lands, with approvals, requested reviewers, and diffs. Protected branches refuse changes that haven't come through an approved merge request, and protection rules can request reviewers automatically (CODEOWNERS-style).
Triggers	Webhooks and the runner	Signed webhooks on commits, merges, and more; the runner runs jobs on a branch when a commit or tag lands.
Proxy / edge server	LAN-first peers and seed nodes	Machines in the same office share data directly; a seed node on your own hardware keeps a warm local copy.

Bringing a depot across

There's no Perforce history importer. Swarmfile can't read a Perforce server's revision history and replay it as commits. What you can bring across is the **current state** of a depot or stream, and that's usually what a team needs to start working.

1. **Get the head revision on disk.** In a Perforce workspace that maps the depot or stream you're moving, `p4 sync` to the revision you want to start from.
2. **Create the Swarmfile project**, and decide its branches (one per stream you're moving, usually starting with `main`).
3. **Import with `swarmfile-migrate`**. Dry-run first to check what would be imported and that your excludes are right:

```
swarmfile-migrate --source /p4/workspaces/game-main \  
  --org-id acme-games --project-id game \  
  --exclude ".p4config" --exclude "**/.p4ignore" \  
  --dry-run  
  
swarmfile-migrate --source /p4/workspaces/game-main \  
  --org-id acme-games --project-id game \  
  --exclude ".p4config" --exclude "**/.p4ignore" \  
  --state-db ./migrate-state.db --verify --report-json ./migrate-report.json
```

`--state-db` makes a large import resumable, `--verify` checks each file after upload, and `--report-json` gives you a record of everything that arrived. The [swarmfile-migrate](#) reference has every flag, and [Migrating Existing Data](#) walks through a full run.

4. **Commit and tag the import** (for example `swarmfile tag create imported-from-p4`) so there's a named starting point.
5. **Recreate permissions and exclusive-checkout habits:** ACLs for restricted folders, and edit locks (or Windows worksharing) for the binary files your team used to `+l` .
6. **Keep Perforce read-only for older history.** Leave the server up in read-only mode (or archive it) for as long as you need to look up who changed what before the move. Swarmfile's history starts at the import.

Moving one stream at a time works fine; Swarmfile doesn't need to own the whole depot on day one.

Where it's genuinely different

- **No sync, no have-list.** You never decide which revision your workspace holds; the drive shows the branch as it is. Large files you haven't opened take no disk space until you do. To keep a set of files local for the road, use [Pack & Go](#).
- **Saves are shared immediately by default.** Perforce shares nothing until submit. Swarmfile's default mode lands each save as it's made; switch a project to staged mode, or open a changelist, when you want Perforce-style batching.
- **History moves forward.** Reverts, restores, and rollbacks are recorded as new commits rather than edits to old ones, and `commit --amend` changes only a message.
- **Whole-file versions.** Like Perforce for binaries, every version is a whole file (stored deduplicated, so unchanged parts of a large file cost nothing to keep).

Not supported

- **No Perforce history importer.** Import the head revision as above and keep the old server for history.
- **No game-engine source-control plugins.** There's no Unreal Editor or Unity source-control provider for Swarmfile, so the editors' built-in check-out and submit buttons don't talk to it. Teams work on the mounted drive directly and use the Desktop App or the `swarmfile` CLI for locks and commits.
- **No Perforce-compatible commands.** The `swarmfile` CLI uses its own verbs (and understands many git habits, see [Coming from Git](#)); it doesn't accept `p4` command lines.

If one of these is a blocker for your move, [tell us](#) — it helps us decide what to build next.

Migrating Existing Data

`swarmfile-migrate` is a standalone CLI for bulk-importing data you already have into a Swarmfile project — a NAS share, a flat list of file paths, or an S3-compatible bucket. Use it the first time you bring an existing archive onto Swarmfile, rather than dragging everything through your mounted drive by hand: it talks to the hub directly rather than through a mount, uploads files concurrently, and resumes from a local state database if the run is interrupted.

Bringing a **git repository** rather than a file tree? Use `swarmfile git import` instead — it carries the commit history, branches and tags. `swarmfile-migrate` moves files, not version-control history.

Figure omitted in print (animation): Importing an existing NAS library: one command reads the share, verifies every file by content, and the library appears as an ordinary drive on another machine. - <https://swarmfile.com/demo/nas-migration.gif>

This page is a practical walkthrough. `swarmfile-migrate` is installed on `PATH` with the desktop app on every platform. For the complete flag reference, see [swarmfile-migrate](#).

A realistic example

Importing a local project directory, skipping cache and scratch files:

```
swarmfile-migrate --source /Volumes/nas/ProjectArchive \  
  --hub-url https://hub.swarmfile.com \  
  --org-id acme-films --project-id feature-01 \  
  --exclude "*.cache" --exclude "**/tmp/**"
```

This walks the source directory, applies the gitignore-style exclude patterns to skip anything you don't want copied, and uploads the rest into the target project.

What to expect while it runs

Migrate reports progress as it goes and retries failed transfers automatically, with backoff, up to 3 attempts by default (`--retry`) — a dropped connection partway through a large file doesn't fail the whole run, it retries that transfer. If the process itself is interrupted (you kill it, the machine sleeps, the network drops entirely), it's safe to just run the same command again — a file that had exhausted its retries in the earlier run is still eligible on the next invocation.

Resume

Migrate keeps a local state database tracking what's already been uploaded. On a second run with the same source and target, it picks up where it left off instead of re-uploading everything from scratch. You don't need a separate `--resume` step or flag — re-running the same invocation is the resume mechanism.

Verification

After each file uploads, migrate checks its content against the source — a CID/hash comparison — so a successful run means the data landed intact, not just that bytes were sent. This is on by default; pass `--verify=false` to skip it if you trust the transport and want to roughly halve I/O on a very large migration. Migrate can also produce a JSON report of the run, which is the artifact to keep if you need to audit exactly what was migrated, skipped, or retried.

Deployment Topologies

Swarmfile is the same product in every deployment below. What changes is where the bytes travel — and that is a setting, not a different install or a migration. You can start central, move to LAN-first when a second person in the same office starts complaining about the same download, and be back again the same day if your security team asks.

Most topologies below are shown as a real recording. A few have nothing to *show* in a terminal — a second site, a headless seed node serving an office — so those are drawn as a diagram instead of faked as a recording.

One thing is true of every topology here, so it is worth watching first: opening a file never downloads more than the bytes your app reads. See [Working with Files](#) for that recording.

Central: hub and spoke

The shape most teams already run. One central store; every machine talks to it and never to each other.

Figure omitted in print (animation): A central share: a file saved on one machine appears on another and both fetched it from the central store — zero machine-to-machine transfers. - <https://swarmfile.com/demo/hub-and-spoke.gif>

Choose this when machine-to-machine traffic on the corporate network is disallowed, when your sites have no meaningful local network between them, or simply because it is the model everyone already understands. It works, and the drive behaves identically.

The cost is at the end of that recording: **zero blocks received from other machines**. Ten editors opening the same 200 GB sequence means ten downloads of it.

LAN-first: one download per office

The same project, with machines allowed to serve each other.

Figure omitted in print (animation): Three machines sharing one drive: a file saved on one appears on the others, and the peer-transfer count proves the bytes came over the LAN. A machine is taken offline and the file stays readable. - <https://swarmfile.com/demo/lan-resilience.gif>

The first person to open a file pulls it from the cloud. Everyone else in that office gets it from them, over the local network, at local-network speed. The recording ends with the same counter the central one does — this time in the dozens.


The second half matters just as much: a machine is taken **offline** and the file is still readable from the others. Erasure coding and peer replication mean the office does not depend on whoever happened to create a file still being at their desk.

LAN-first also applies *before* a file reaches the cloud at all. When someone asks for part of a file that is still uploading — `swarmfile fetch`, or a follow job tracking a viewer — the uploading machine hands those blocks straight off its own disk to a peer in the same office, rather than everyone waiting for them to climb the office uplink and come back down from storage. That needs the project's blocks to be peer-servable (an encrypted project, or a public one); blocks that don't qualify wait for the cloud. Serving a peer straight off the uploader's disk also only happens on a LAN, deliberately: over the internet it would put every block on the same congested uplink twice, once to the peer and once to the cloud, which is slower for everybody. It is also the reason this workflow is far more comfortable in one building than across two — on a LAN the constraint is the uploader's own disk, not its uplink.

Choose this when several people work in the same building on the same material. It is the default for a reason.

Multiple sites

A studio with people in two buildings, or two cities, tells each machine which office it is in. Machines then **prefer** peers in their own office: same-site peers are offered first, then any always-on seed nodes, then whatever else is around.



Preference rather than a hard boundary is the useful behavior. A colleague two desks away is the obvious place to get a file from, and that is what happens in practice. But if your site is quiet — one person in early, everyone else still asleep in another timezone — a machine is not forbidden from using a distant peer that happens to be reachable, and it falls back to the cloud when nothing better is available. You get locality where locality exists, without a rule that strands someone working alone.

Nothing about this needs configuring per project. Sites differ; the project, its history and its permissions do not. Add a second person at a quiet site and that site starts sharing locally on its own.

One person, working remotely

A single editor at home has no peers by definition, and the drive behaves like the central topology above — everything comes from the cloud.

Which raises the obvious question: what happens when the cloud isn't there? A file you've been working with stays fully readable, because it's already on your machine — and a file you never opened fails at once with a clear answer rather than hanging your application, because there's no connection to fetch it over. When the connection returns, everything is simply there again.

Figure omitted in print (animation): The connection to the coordination service drops: a file already opened still reads in full from the local machine, while one never downloaded fails immediately instead of hanging; reconnecting restores everything. - <https://swarmfile.com/demo/offline-resilience.gif>

One thing does change underneath: erasure coding exists to mask WAN latency, and it switches itself off once a machine can see two or more peers on its local network. A solo remote machine keeps it on and a busy office machine does not, without anyone choosing. If you are evaluating Swarmfile with one person on a home connection, that is the configuration you are testing.

Seeding from what you already have


Neither topology is any use if getting your existing library in means a weekend of copying.

Figure omitted in print (animation): Importing an existing NAS library: one command reads the share, verifies every file by content, and the library appears as an ordinary drive on another machine. - <https://swarmfile.com/demo/nas-migration.gif>

The import reads your NAS and never writes to it, so you can point it at a live share. It is resumable — interrupt it and run it again — and every file is verified by content, not by size and timestamp. Full detail in [Migrating Existing Data](#).

Keeping the NAS as a warm tier

Importing does not mean retiring the hardware. A machine you own — the NAS, an old workstation, a dedicated box — can join permanently as a **seed node**: it fetches the whole project, mirrors every block to cloud storage, and keeps a warm local copy of the working set to serve everyone in the office over the LAN.



The difference from an ordinary machine is that a seed node is always there. LAN-first already means the first person to open a file supplies everyone else, but that depends on that person being at their desk with the machine awake. A seed node removes the dependency: its cache always holds the recent working set, and older blocks re-fetch on demand (the whole project is mirrored to cloud storage, so nothing is ever unavailable).

Seed nodes are a **Pro plan feature**. This is enforced rather than advertised — a seed node on a plan without the entitlement is refused by the service, and `swarmfile doctor` reports it plainly rather than leaving a repeating warning in a log nobody reads. Setup, sizing and hardware guidance is in [Self-Hosted Seed Nodes](#).

The same shape works in a data center rather than an office, for a team with no single main site: an always-on warm tier that belongs to nobody in particular.

Larger offices: deciding where the shards live

In an office of a handful of machines, whoever opens a file ends up holding it, and that is enough. Past a certain size, leaving your fault tolerance to whoever happened to open what is not a plan.

Figure omitted in print (animation): Deliberate shard placement in a five-machine office: turning it on assigns each piece of a file across three of the five machines, and swarmfile shards names exactly which — a decision, not chance. - <https://swarmfile.com/demo/shard-placement.gif>


Turning on deliberate shard placement makes it a decision instead: each piece of a file is assigned across the office's machines at a replication factor of three, so any given file survives losing machines rather than surviving because someone happened to have opened it. `swarmfile shards <path>` shows the live assignment for any file — which machines should hold each piece, and which ones actually do.

Worth turning on when an office is large enough that "who has a copy?" stops being a question you can answer by asking around.

One caveat on corporate networks: deliberate shard placement needs a roster of LAN peers to spread across, and LAN peers are normally discovered by mDNS multicast. On segmented networks where mDNS multicast is blocked, set `lan_from_office` (`SWARMFILE_LAN_FROM_OFFICE`, or the config-file key — see [Engine Config File](#)) so same-office peers count as LAN. Both the `office_id` label and this flag apply live via `swarmfile office set <office_id> [--lan-from-office]` within one peer-discovery tick (~60s, no restart) — unlike enabling shard placement itself, which needs a restart.

Mixed Windows and macOS offices

A normal post or design floor is not one platform, and Swarmfile does not ask it to be. Windows machines mount through WinFsp and get a drive letter; macOS mounts through FUSE and gets a volume. Both are ordinary local paths to the applications on top, both are peers to each other, and a file written on one is a file on the other.



One practical consequence worth knowing up front: **the drive's location is decided per machine, at runtime**. On Windows, a taken drive letter is only silently relocated to a free one when it was never an explicit choice — still sitting on the app's own default or a letter it auto-picked — and a free candidate actually exists. If you (or a script) explicitly set the letter, Swarmfile honors that choice and reports the conflict by naming the occupant instead of quietly moving you off it; likewise if every candidate letter is also taken, the mount fails and names the occupant rather than relocating. So no script, template or onboarding document should hard-code a path — read it from the Desktop App's Status pill, or from the app, which knows.

Render farms and headless machines

Not every machine that needs the project has somebody sitting at it. Render nodes, watch folders, transcode boxes and QC stations are usually many readers and few writers, which is the shape LAN-first is best at: the first node to pull a frame sequence supplies the rest of the farm locally instead of every node fetching it independently.

Two things make this work well in practice. Mount the farm **read-only** where it only needs to read, so an automation cannot write into the project by accident. And give the farm a **seed node** — a machine that keeps the project warm and mirrored — so the first read of the day is a local read too, rather than depending on whichever workstation happens to be awake.

If the farm itself is cloud-hosted — a cluster of render nodes in your own cloud account rather than machines in a physical office — there's no single LAN to put that one seed node on. The farm can still get redundancy among its own nodes instead of leaning on a single point of failure: enable seeding and set the same `office_id` on each node (no shared network required — see [Self-Hosted Seed Nodes](#)). Each seed node independently fetches every block and mirrors the full project to cloud storage, so the project's cloud copy is complete on its own — a stronger guarantee than deliberate shard placement's replication factor of 3 — and there's no extra step needed on top of it for an all-seed-node farm.

The other thing an unattended machine usually needs is to be pinned to an exact version, not "whatever main happens to be right now" — a render should use the frames as they were when the job was submitted, even if someone pushes a fix five minutes later. Tag the commit a job was submitted against, then have each node branch from that tag: one command, no Desktop App or browser interaction, and the node's copy is frozen to precisely that content regardless of what lands on main afterward.

Figure omitted in print (animation): A headless render node signs in with an API key, tags the submitted version, forks a branch pinned to that tag, and hydrates the scope resident before the job. An artist then commits a new scene to main — and the render node, on its pinned branch, still shows only the frames that were submitted. - <https://swarmfile.com/demo/render-farm.gif>

```
swarmfile commit -m "submit render job 482"
swarmfile tag create render-482
swarmfile branch create farm-482 --from-tag render-482
swarmfile branch switch farm-482
```

`tag create` with no `--at-seq` tags the branch's current head — exactly the commit the submission script just made, with no separate step to look up its commit number.

Mounting sparse is right for interactive work, but a render node has no interactivity to hide latency behind: the first frame it opens still has to fetch before the job can proceed, and if that node has no LAN peer or seed node nearby, every frame is a cloud round trip on the critical path. `swarmfile hydrate start` forces a scope fully resident up front instead — download and pin every file under a path (or the whole project as this mount's branch sees it) before the job starts, rather than fetching lazily as frames are opened:

```
swarmfile hydrate start ./Projects/shots/seqA --yes
```

Hydrated files are pinned against the cache's normal eviction, so they stay resident until you explicitly `swarmfile hydrate release` them (or delete the local copies outright with `swarmfile hydrate free-space` once the job is done) — useful for a node that needs to keep working with no network at all for the length of a job. Pair it with the branch-from-tag pin above: hydrate the exact scope a job was submitted against, and the node has everything it needs, locally, before a single frame renders.

A farm node authenticates with a project-scoped **API key**, not a human session. Mint one from the web dashboard, the Desktop App, or the CLI — an admin or owner picks a name and a single project, and gets back a secret shown exactly once:


```
swarmfile api-key create farm-node-12 --project proj_9f2a
```

Set it as `SWARMFILE_API_KEY` on the node (alongside `SWARMFILE_CONFIG` pointing at a per-project config file) and the engine skips sign-in entirely — no browser, no refresh token, no impersonating whoever happened to mint it. The key is walled off to that one project: it can't read or write anything else in the org, even projects the admin who minted it can see, and revoking it (`swarmfile api-key revoke <id>`, or from the dashboard/Desktop App) kills that one node immediately without touching anyone's own session. See [CLI: swarmfile](#) for the full `branch / tag / hydrate / api-key` reference.

For the coding-agent variant of this shape — several agents, each with its own labeled mount and branch, sharing one dependency cache — see [Coding Agents](#).

What runs where — and what you can host yourself

Worth being direct about, because it comes up in every security review.



On your hardware: the file data. [Dedicated storage](#) is the default for new paid orgs — an isolated bucket of their own rather than a shared one — and on Enterprise [bring-your-own-storage](#) can make an S3-compatible bucket on your own account the org's primary block storage. Seed nodes keep the project warm locally (their cache budget bounds the disk; every block is mirrored to your storage), and LAN-first means bytes move machine-to-machine without a round trip to anyone's cloud. A studio with a seed node in the building serves most reads from its own hardware.

Where it lives: every paid plan can pin an organization's metadata and block storage to a real EU or US jurisdiction at signup, at no extra cost (Free uses shared storage and isn't eligible), and an individual project can carry its own [residency pin](#) independent of the org's. Enterprise can additionally provision [FedRAMP / FedRAMP-High storage regions](#) through sales (storage only — Swarmfile is not FedRAMP-authorized).

Managed by us, by default: the control plane — metadata, permissions, identity, and the coordination that lets machines find each other. Most customers run this as a hosted service.

Self-hostable on Enterprise: the control plane can also run entirely on infrastructure you control, on a self-hostable runtime compatible with the platform it's built on — the same code, not a fork or a lesser rewrite. This is available on Enterprise today; [talk to us](#) about what it would involve for your environment. A fully air-gapped deployment, with no outside connectivity at all, is not available today.

So Swarmfile suits an organization that wants its **bulk data** on infrastructure it controls, even on the standard hosted plan — and if you need the **coordination layer** inside your own perimeter too, that's an Enterprise option rather than something the rest of this page's topologies can get you. Raise it early in a procurement conversation so it gets scoped correctly.

Encryption is the thing worth reading next if this section is what you came for: per-project keys, with an end-to-end tier where the service cannot read your content at all. See [Security](#).

Guests, clients and consultants

Not everyone who needs a file needs a machine in your topology. External collaborators get a named, revocable read — or read-and-comment — view of one folder, with no paid seat and a full audit trail — a colorist, a client reviewing a cut, a consulting engineer. Covered in [Sharing & Collaboration](#).

Choosing

If this is true	Use
Machine-to-machine traffic is not permitted on your network	Central, enforced org-wide
Your people are spread across homes and cities, rarely co-located	Central
Several people share a building and open the same large files	LAN-first
You have multiple offices	LAN-first at each; they are independent
You have an existing NAS you want to keep using	LAN-first plus a seed node (Pro)
Nobody has a main office, but you want a warm copy somewhere	A seed node in a datacenter (Pro)
An office is big enough that "who has a copy?" is unanswerable	LAN-first plus deliberate shard placement
Someone outside the company needs to see one folder	An external collaborator, not a topology change
Your floor is a mix of Windows and macOS	Either topology; both mount natively
Many machines read, few write (a render farm)	LAN-first, mounted read-only, plus a seed node
Everything including the control plane must be inside your perimeter	Enterprise self-hosted control plane (fully air-gapped is not available today) — see "What runs where" above

Switching is a setting, not a migration: the files, history, and permissions are unaffected. Try one, measure it, change your mind.

Two things worth knowing about who holds that setting. Each machine can choose for itself — but an organization can also set its data-plane policy to **hub-only for everyone**, and that decision overrides the local setting on every machine, whatever it is set to. If your security position is "no machine-to-machine traffic on our network", that is enforceable centrally rather than left to each person to configure correctly.

The reverse is not true: an org that permits LAN-first does not compel it. A machine on a network where peer traffic is blocked simply finds no peers and falls back to the cloud, which is the central topology arriving at itself.

Working with Files

Your mounted drive is a normal filesystem to any application — Resolve, Premiere, Revit, QGIS, Explorer, Finder, whatever you point at it. This page covers what's different about working with files on Swarmfile: the web dashboard's view into the same data, and the mechanics of saving, presence, and locking that run underneath.

Browsing

The web dashboard's Files view is a file explorer: browse the project's folders and files in the left pane — each row shows its size and how recently it changed — and open one to see an inline preview (images, PDFs, code, Markdown, and generated thumbnails for other types) with its metadata alongside. This is a separate view onto the same underlying storage as your mounted drive — useful for checking file history or activity without opening a native app. On the mount itself, browsing is just normal folder navigation, and it costs no disk space: listing a directory doesn't download anything in it.

When someone else saves a file, it appears in your file manager within a couple of seconds. Your operating system caches directory listings for a short window, so a folder you already have open can be a moment behind; reopening it or pressing refresh always shows the current state.

A window that already has the file open keeps the version it opened — each open handle is served one consistent version, so a teammate's save never swaps bytes under a running app. Reopen the file to see the new version; the Desktop App posts a one-time notice ("... was updated — reopen it to see the new version") when one of your open files has been superseded. Fresh opens and the web dashboard see the new version immediately.

A single folder holding tens or hundreds of thousands of files — a point-cloud tile set, a scan-output dump — is loaded on demand: the web dashboard's Files view fetches a folder's contents when you expand it and renders at most a couple of hundred rows at a time (with a "Show all" reveal), and the Desktop App's file picker only renders the rows currently in view, so both stay responsive on a folder that size instead of the browser or app choking on it.

The web dashboard additionally renders thumbnails and inline previews — including video posters, scrubbable proxies, and point-cloud previews — so you can eyeball a folder without opening a native app. See [File Previews](#).

Opening and saving

Opening a file streams the bytes your application reads — and only those bytes. Those bytes come from the nearest place that has them: the local cache first, then a teammate or another of your machines on the same LAN, then a self-hosted seed node if your org runs one, then peers elsewhere on the swarm, and Swarmfile cloud storage last. WAN-heavy reads can also use adaptive Reed-Solomon erasure coding (10+4) so a few slow peers don't stall the transfer.

Figure omitted in print (animation): A machine lists a 64 MB file it has never fetched and its local cache does not grow at all; it then reads one megabyte from the middle and the cache grows by about one megabyte, not sixty-four. - <https://swarmfile.com/demo/sparse-streaming.gif>

The recording above is the whole idea in three commands: listing a file costs nothing, and reading part of one costs that part. Scale it up and it is why a 200 GB sequence opens on a laptop with 40 GB free. Saving is asynchronous: a write returns as soon as it's queued, not once it's fully uploaded. A durable local upload queue handles the actual transfer in the background, with automatic retry on failure.

Figure omitted in print (animation): A 48 MB save returns at local-disk speed while dozens of blocks are still queued for upload, and the queue then drains to zero on its own. - <https://swarmfile.com/demo/zero-stall-writes.gif>

That recording is the same point from the writing side: the save returns at local-disk speed with dozens of blocks still queued behind it. Because the queue is durable, it survives a restart or a dropped network connection — pending uploads resume automatically rather than being lost. You are never blocked waiting for a save to finish uploading before you can keep working.

Reading the sync status

The Desktop App's header keeps one line of truth about that queue. While work is in flight it says **Syncing N** — N counts the work still on its way to the hub, **one per save** (a 40 GB save counts once, not once per block): saves whose bytes are still uploading, new files whose creation hasn't been confirmed yet, and saves whose bytes are already up but whose commit hasn't landed. **Synced** means the queue is empty, not merely that a transfer looks quiet: a save whose commit is still being retried keeps the header at *Syncing* until it lands. A busy hub or a sign-in that has lapsed only pauses the queue; neither uses up a save's retries, so a save doesn't give up while you sign back in.

Three of those states are worth knowing by name:

- **Waiting on plan** — everything left is a save the hub is refusing on plan grounds, almost always storage (the org's allowance is full) or the org's spend cap. The retries are automatic, but nothing moves until the limit is raised; a quota/billing-block notification tells you which limit, and the plan panel shows the allowance. An owner or admin raises the spend cap under **Billing → Spend cap**.
- **N changes not syncing** — the hub keeps refusing these saves for another reason, so retrying isn't making progress. They're safe on this computer and Swarmfile keeps trying; the attention indicator points to the **Status pill** in the header, which opens the panel showing the hub's reason. If the hub refused a save because part of its content never arrived, Swarmfile re-sends that content from this computer on its own before asking you for anything. The Status panel lists each file that is still stuck — why, and for how long — with **Discard this change...**, which drops the unsynced change and puts the file back to its cloud version. That can't be undone, so it asks first. From a terminal: `swarmfile sync stuck` and `swarmfile sync discard <entry-id>`.
- **N changes failed** — some work stopped retrying on its own. The header's attention indicator names each one, and the banner above the file list offers **Retry**; an upload that failed can also simply be saved again. On Windows, saving the file again also clears the error badge Explorer shows on it.

The header counts the project you're looking at. Saves you made in another project keep uploading in the background after you switch away from it, and the Status pill adds **N saving in other projects** (or names the project when only one has saves left) until they land. This project can read **Synced** while that's true. A change in another project that the hub is refusing is named there too, and stays on this computer until it's resolved in that project. `swarmfile status` prints the same: its pending uploads line adds "(N more in other projects)", followed by one line per project with its saves still uploading or held.

A file you create is listed in the Desktop App immediately, before the hub has confirmed it: the row reads **Syncing — this one isn't on the hub yet** until the create lands (seconds on a healthy connection), and you can already rename or delete it from the app — applied on this computer, not waiting for the hub. A create the hub refuses but *holds* — root-create authority on a protected project, say — names the reason on the row instead and keeps retrying (see [Permissions](#)). A teammate's file

that was just created shows **Waiting for the hub** until the hub confirms it (it can still be retracted by whoever created it). Moving or copying a file whose create hasn't landed, and moving, copying, uploading or importing anything *into* a not-yet-confirmed folder, wait for the hub.

While a large file is still uploading

The queue draining in the background is invisible on a fast link. On a slow one it isn't: a very large file on a modest uplink can take hours or days to finish. Until it does, everyone else keeps reading the last committed version — the in-flight save never replaces it. (A file's *first* upload is the exception: it can be opened while it is still arriving — see [Opening an in-progress upload directly](#).)

While that transfer is running, the file's page in the web dashboard shows an **Uploading** badge, along with which machine is sending it. The row itself still describes the version everyone can currently read — same size, same date, and Download still works — because a version in progress never replaces the one people are already using. Nothing becomes unavailable because a colleague started saving.

If you need part of that new version before it has all arrived, you can ask for it, and the uploading machine will send that part first:

```
swarmfile fetch ./Projects/reel-04.r3d --tail 209715200
```

That is the last 200 MB of a file still in transit. The request travels to whichever machine is uploading, that machine moves the blocks covering your range to the front of its queue, and you get them in roughly the time it takes to send *that range* — rather than waiting out the whole file. Uploads are sent in order from the start of the file, so without asking, the end is the last thing to arrive.

It is worth knowing what this does and doesn't do. It changes **which bytes arrive first**, not how fast the link is: the last 20% of a 5 TB file is still 1 TB, and 1 TB still takes as long as 1 TB takes. What it removes is the wait for the other 4 TB you didn't need.

The bytes land in your local cache, pinned so they aren't evicted before you open the file. Release them with `swarmfile hydrate release` when you're done — see [the CLI reference](#) — or delete them outright with [Free up space](#). Reading the file normally through the mount is unaffected and gives you the committed version — unless the mount has been opted into streaming, below.

Watching a file as it uploads

A single range is the right shape for "give me the last 200 MB". It is the wrong shape for watching, because a viewer moves: it plays forward, and it seeks. `--follow` tracks that.

```
# Start following from the beginning
swarmfile fetch ./Projects/reel-04.r3d --follow --version pending

# ...and tell it where the viewer has got to
swarmfile fetch-seek 5000000
swarmfile fetch-status
```

Instead of completing at a fixed range, a follow job keeps a buffer ahead of wherever the reader is and re-targets when it moves. The uploading machine is told about the new position immediately, and — this is the part that matters — the new request **replaces** the old one rather than queueing behind it, so it stops working on the stretch you scrubbed away from. A follow job ends on its own when it reaches the end of the file, or after five minutes with no `fetch-seek` — a status that stops advancing means it finished, not that it broke.

How far ahead it buffers is measured, not configured: the read position's own rate of advance *is* the consumption rate, so a 6 Mbit/s proxy gets a small window and a 200 Mbit/s master gets a proportionally larger one, both covering the same number of seconds. `SWARMFILE_STREAM_BUFFER_SECS` sets that number (default 30). A seek is not mistaken for playback.

Opening an in-progress upload directly

Everything above puts bytes in the cache. There is also a mode where a file that has never finished uploading is simply **openable** — it appears at its full eventual size and reads work, waiting briefly where the bytes have not landed yet.

This is always on. A read into a range that hasn't arrived has to wait, and a filesystem operation that waits too long doesn't stall one read — it takes the whole drive with it. The wait here is strictly bounded (15 seconds, see `SWARMFILE_STREAM_READ_WAIT_SECS`) and a read that outruns the upload returns "try again" rather than an I/O error, precisely so an application doesn't conclude the file is damaged. See [the config reference](#) for the full behavior.

Two things to know about the scope:

- It applies only to files with **no** committed version yet — a new master landing for the first time. A file you are already reading never changes size or content underneath you because somebody started a new save. That is the same rule the rest of the product follows, and it matters more than the feature does.
- If the uploading machine is on your **LAN**, it will serve those blocks to you directly, from its own disk, before they have reached the cloud at all. Across the internet it won't: that would put every block on the same congested uplink twice. Direct serving also requires the project to be **encrypted or public** — a *private* project's plaintext blocks are never handed out peer-to-peer; a public project's cleartext blocks are already anonymous-readable, so they are. Where blocks aren't peer-servable, in-flight reads wait for the cloud like any other.

The real limit. None of this makes the link faster. It makes the right bytes arrive first. A 4 TB file on a 100 Mbit/s uplink delivers about 12 MB/s however well it is ordered — comfortable for a proxy, not for a full-resolution master. The reliable version of this workflow is "watch the proxy while the master uploads", and on a LAN, where the uploader's own disk is the limit rather than its uplink, considerably more than that.

Freeing up space

Everything you open stays in a local cache so the next read is instant, and the cache trims itself when it reaches its size limit. The *first* read of a cloud-only file is the one exception: the engine fetches on demand and reads ahead as playback advances, so a large media file can take a few seconds to settle while it builds that lead — and it only ever pulls the byte ranges you touch, never the whole file. Once those bytes are cached, the file is local-fast like any other.

To make room on purpose — or to hand a laptop's disk back before traveling — use **Free up space**. It deletes the downloaded copy so the file is cloud-only on this computer; the file stays listed and downloads again the moment something opens it. The file list shows that state beside the size: a quiet **cloud** tag (or **partly local** once some of it is cached) marks rows whose full size is on the hub, not on the disk.

- **A file or folder:** in the Desktop App, open the row's menu and choose **Free up space**. On Windows, you can also right-click it in Explorer and choose **Swarmfile** → **Free up space**.
- **The whole project:** in the Desktop App, open ⋮ **Project actions** on the active project in the left rail and choose **Free up space (whole project)**.... It shows roughly how much it will free before you confirm, and it also clears older versions of files that are still cached.
- **From a terminal:** `swarmfile hydrate free-space [path] ; add --dry-run` to see what it would free without deleting anything.

It is safe to run at any time:

- **Changes that haven't synced yet are never touched.** A file with a save still on its way to the cloud is skipped, and the result says how many were kept.
- **Make available offline is turned off** for what you free up, since keeping it offline and freeing its space contradict each other.
- **Content something else is holding stays**, such as work staged in a changelist, a Pack & Go reservation, or a file that is open right now. The result counts these as still partly on this computer.

Executable files

A file can be marked executable: a script, a build tool, anything you run with `./name`. On macOS and Linux, `chmod +x` on the drive sets it and `chmod -x` clears it; the drive shows executable files as `0755` and other files as `0644`. Windows has no executable bit, so edits there leave the flag as it was. From any platform, `swarmfile chmod +x <path>` (or `-x`) sets it. The flag syncs to everyone like any other change, works offline (it's sent when you reconnect), survives safe-saves that replace the file, and becomes mode `100755` in a [git clone](#). On macOS another machine's change can take a couple of seconds to show, the length of the system's file-attribute cache.

Windows alternate data streams

On Windows, NTFS alternate data streams (`file.ext:streamname`) are supported on the mount: applications can create, read, write, enumerate and delete them, a stream's content travels the same upload and commit pipeline as file content and syncs to other clients, and a stream survives a copy or a rename of the base file. The one exception is `Zone.Identifier` — Windows' Mark of the Web — which stays on the machine that wrote it and never syncs (see [Sync Exclusions](#)). Renaming a stream on its own is refused, and Windows extended attributes are a separate mechanism and aren't implemented. See [Filesystem Compatibility](#) for the verified rows and remaining gaps.

Deleting

Delete from the drive exactly as you would any other file — your file manager, or `rm`. Deleted files go to Trash rather than disappearing, and are restorable for as long as your organization's retention window allows: see [Trash, History & Rollback](#).

Two behaviors worth knowing. On the mounted drive a folder is removed only once it's empty: your file manager, `rm -r` and `rmdir /s` empty it for you, one file at a time (PowerShell's `Remove-Item -Recurse` is a known non-starter through the mount — use `rmdir /s` or `Directory.Delete` there), while a plain non-recursive delete of a folder that still has files is refused, as on any disk. In the dashboard, deleting a folder removes its contents with it in one step (they all go to Trash together), so the confirmation there names what's inside; a very large folder still leaves the list at once, the delete showing progress in the toast while its contents follow as a background job. Restoring a folder you deleted from the drive offers its files back too; see [Trash, History & Rollback](#). And if a file is open or locked by someone else, the delete is refused rather than queued: your file manager reports it, and the Desktop App's Presence view shows who has it.

Applied immediately, confirmed in the background

Renaming or moving a confirmed file, or deleting a file or empty folder, takes effect on your machine at once, so a slow or briefly unreachable link can't stall it. The change is sent in the background and shown as a lightweight overlay until the hub confirms it, then re-verified and refreshed in place. Another machine that has already received the change shows the new name right away too, with that row marked **Waiting for the hub** instead of **Synced** until the confirmation lands; one that hasn't received it yet keeps the old name until then. Unconfirmed work is held and retried rather than dropped; if the hub refuses it (for example, you don't have write access), it's retracted and the entry returns. A refused save-by-replace (an app writing a temporary file and renaming it over the original) brings the original back with your unsynced edits to it intact and readable at once. The dashboard's history, Trash, merge-request and RFI views catch up as it lands. The behavior is on by default; an operator can turn it off with `deferred_namespace: false` in the [engine config](#) (`SWARMFIL_DEFERRED_RENAME=0`), which makes these operations wait for the hub again.

Saves project the same way. Once a local flush commits, the update reaches other machines by gossip before the hub round-trip confirms it, and a teammate's row shows **Waiting for the hub** — with the size you saved in its tooltip — as soon as the projection arrives. The file's *content* is the exception: a projected version is display-only, so every machine keeps reading the version the hub has confirmed until the confirmation lands, and an abandoned save is retracted rather than served. Queued saves refresh on a 60-second loop. Size and modification-time display still come from the hub. A machine can be told to fetch a projected version early: Settings → **Prefetch projected updates** (off by default) warms the new version in the background so it opens immediately once the hub confirms it.

Presence

The dashboard shows a real-time "who's editing what" indicator. This is tied to the entry lock lifecycle, not a free-running heartbeat — it reflects files that are actually open and locked, not just which teammates are online. If presence shows someone on a file, they hold (or recently held) a lock on it. This is the authoritative "who has this file" signal: it means a real Swarmfile lock exists, and writing is genuinely blocked while it does. A teammate appears when they first change a file and drops off when they close it — an app that opens a file read-write only to show it (QuickTime Player opens every movie that way) takes no lock and never appears. A dashboard action that locks a file only briefly, such as rolling back a version, never shows as editing. When a teammate is creating many files at once, the Files view also shows "N files arriving from ..." until the hub has recorded them all. In the Desktop App, a presence row that arrived as a LAN hint from a teammate's engine rather than in the latest hub snapshot carries a muted `LAN` marker, so a hint-derived "editing" state is visibly distinct from hub truth. Presence is only in the signed-in dashboard, never on public project pages or share links.

Separate from locks, a teammate can opt in to share what they are viewing. With Settings → **Share what I'm viewing** on (off by default, per machine), a file they have open right now shows a muted **viewing** marker on your rows and in the Presence view — alongside, not instead of, the lock-derived editing state, and it never blocks anyone. Viewing hints travel peer-to-peer only: the hub never stores them, they expire on their own (after roughly 90 seconds), and they clear as soon as the teammate moves on. Several viewers of one file roll up into one badge with their names in its tooltip.

"In use" signal (Windows)

There's a second, softer indicator that answers a related but different question: not "who holds a lock" but "is some Windows application sitting on this file right now". Many desktop apps — Excel, and a number of CAD tools — open a file for writing and deny write-sharing to everyone else while they have it open, without ever taking a Swarmfile lock. When the Windows drive sees an open like that, it reports it best-effort so teammates get a heads-up before they start editing the same file.

Where it shows up: an informational **i** **in use** pill in the web dashboard's Files table, and in the Desktop App's file list, naming who has the file open and on which machine. It's `--info`-toned on purpose — a note, not an alarm.

Two things to be clear about:

- **The pill is a heads-up; the enforcement is separate.** The pill itself doesn't block anything — it's a best-effort display. But the *open* it reflects (write intent with write-sharing denied) is exactly what [worksharing](#) enforces across machines: on Windows, a second machine that opens the file the same way is refused, so two people generally can't open the same central model for editing at once. The pill tells you; worksharing stops you. To reserve a file you aren't actively holding open in an app, use an [entry lock](#), below.
- **It is distinct from Presence.** Presence (above) is driven by Swarmfile's own explicit lock lifecycle. The "in use" pill is a passive read of what a Windows app is doing to the file — it has no lock lifecycle of its own; the cross-machine blocking rides on the underlying open (see worksharing, above). The two can disagree (a file can show "in use" without an explicit entry lock, and vice versa), and that's expected.

Because it's best-effort and poll-driven, it's not instantaneous: the dashboard and Desktop App refresh it about once a minute, and a file that stays open unusually long can eventually drop off the signal even while it's still open. Treat it as a hint, not a source of truth. It's Windows-only — other platforms don't surface this open mode.

Entry locks

Locking a file blocks other users from writing to it until you release the lock or it expires. There are two kinds, and they behave differently:

- **The automatic entry lock** a file takes from the moment your application first writes to it until it closes the file. Opening a file read-write does not take it on its own; the first write, truncate, or overwrite does, and if someone else holds the file that write is refused ("try again") rather than the open. On Windows, opening a file for writing while someone on another computer is editing it is refused at the open instead, as "file in use", which apps report the way they report any file open elsewhere. It has a **60-second TTL** and is renewed with a heartbeat while the file stays open — you don't need to manually extend it. If your app crashes or your machine loses connectivity, this lock expires on its own within 60 seconds rather than staying stuck. If the lock is lost while the file is still open (it expired, or someone else took it), macOS and Linux refuse the next save — the app sees a generic `EIO` ("Input/output error"); the Desktop App names the real reason. Windows keeps the save on your computer instead. It syncs normally if the lock can be taken again; if the file was changed elsewhere in the meantime, your save becomes a [conflict](#) to resolve rather than overwriting the other version. Two saves racing on the same file with no lock held never silently overwrite each other either: the losing version is kept — as a parked conflict to resolve, or as a copy beside the file (`report (conflict-1a2b3c4d).pdf` ; a dotfile or extension-less name just gets the suffix) — so no edit is dropped. (The deferred-namespace opt-out changes *when* the racing save is applied — synchronously instead of through the queue. In that mode a safe-save aside still lands its copy, but a racing rename surfaces as a save error instead of a copy, and the next attempt tries again — nothing is ever dropped silently.) Until it does one or the other, the Desktop App's file list marks the file as kept on this computer, and counts it under the files that need attention.
- **An explicit edit lock** you take yourself. `swarmfile lock <path>` is long-lived — 7 days by default, up to 30 — and its lease is renewed while you hold it. A `git lfs lock` takes a fixed 30-day lease that nothing renews, so a lapsed one releases itself; take it again if the work outlives it. Both last until you release them or the lease lapses; see [Entry locks in the CLI reference](#).

A bulk [Pack & Go reservation](#) is a different mechanism: it reserves a whole folder or project for offline work rather than locking one file, though it blocks writes the same way while it's active.

The lock is keyed to the **user and machine**, not to the mount: two drives of the same engine on one computer share that identity, and share the lock: closing the file in one drive doesn't release it while another drive still has it open. An entry lock taken through one drive does not by itself keep a sibling drive out. When several drives on one machine must not edit the same file — a fleet of agents, say — open them with [exclusive-write mode](#).

The web dashboard's project-scoped **Locks** tab (admins and owners only) lists every file currently locked in the project — standard and 30-day edit locks alike — with who holds each one, and lets an admin release their own lock from there directly. A plain member releases their own edit lock the same way they took it — close the file (the short-lived automatic kind clears on its own) or `swarmfile unlock <path>` for an explicit edit lock — rather than through the dashboard. Releasing someone **else's** lock, admin or not, still goes through the request-unlock flow below.

Asking someone to release a lock

If you need a file someone else is holding open, right-click it in the web dashboard's Files view and choose **Request unlock** (also reachable from the Windows Explorer `Swarmfile` ► submenu on a mounted drive, or from the CLI with `swarmfile unlock-request create <path>`). The current holder gets a notification — [in-app only, not emailed](#) — with a **Grant/Deny** right on the notification row, whether that's a modal in the Desktop App naming who's asking or the web dashboard's Inbox (or, from a script, `swarmfile unlock-request list / respond` to see and answer pending requests headlessly). They can release the lock from there, or ignore it and let it expire on its own the normal way — a short-lived entry lock clears when they close the file or within 60 seconds if their app or machine has gone away; an explicit edit lock runs out at the end of its lease (up to the 30-day maximum). There's no force-release button anywhere — asking is the mechanism, full stop; a still-pending request you made can be withdrawn from the same file's right-click menu (**Cancel unlock request**) or with `swarmfile unlock-request cancel <request-id>`. The one exception is the `git-LFS` surface: an entry lock and a `git lfs lock` on the same file are the **same** server-side lock, so a `git lfs unlock --force` can break another user's lock (git-LFS's own `--force`, not a drive/dashboard action).

Byte-range locking

Entry locks cover a whole file. For native worksharing apps that need to lock just part of a file — the way a Revit-class BIM tool locks the elements one person is editing without blocking everyone else out of the same model — Swarmfile exposes a separate, lower-level byte-range locking API. This is what a CAD/BIM plugin

integration would call directly. The API is enabled only on projects set to the Revit-worksharing project type — at creation, or later from **Project settings** or `swarmfile project set-type revit_worksharing` (see [Worksharing](#)); the engine runs the byte-range lock manager only there, and other projects refuse those requests. The [swarmfile-brlock](#) CLI is the reference client for the byte-range lock API.

File and folder watch

Watch is distinct from the general "someone else changed a file" bell notification you see for activity across a project. Watch lets you subscribe to a specific file or an entire folder and get notified when it changes — useful for scripting a rebuild step or keeping an external tool in sync with a subtree of a project without polling it.

If the drive stops responding

Occasionally the drive can stop answering: files won't open, a Finder or Explorer window sits there, and the application waiting on it may not even force-quit cleanly.

That behavior is deliberate, favoring write safety. The drive is mounted to *wait* rather than to fail, because a filesystem that returns an error partway through a write can leave an application believing a save succeeded when it didn't — which on a multi-hundred-gigabyte project file means silent corruption. The cost of that choice is that a drive which stops answering waits indefinitely instead of erroring.

Swarmfile watches for this and reconnects the drive on its own where it can. While that's happening the Desktop App reports the drive as unavailable rather than continuing to show a healthy mount; when it reports the drive as available, it is.

When it doesn't clear by itself, open the Desktop App, go to Settings → Diagnostics, and use **Repair drive** (macOS and Windows; Linux has no in-place repair, so restart the engine or remount instead). It reconnects the drive in a few seconds and touches nothing else — your queued uploads are held in a durable local queue, so nothing waiting to upload is lost, and you don't need to quit and reopen Swarmfile.

If the drive keeps dropping instead of recovering — the Desktop App's "keeps disconnecting" card — it is usually a leftover mount from a killed session. The engine clears those automatically before each mount, so **Repair drive** normally reconnects it; if Diagnostics still reports a stuck mount, only a computer restart clears it (the state lives in the kernel, not the install), and **Reinstall...** is for a genuinely broken install. For working out whether the underlying problem is your network, your account, or the service, run [swarmfile doctor](#) — its output is what support will ask for first.

Working Offline (Pack & Go)

Most of the time your drive is a live view of the project: files stream on demand and saves sync in the background. But sometimes you need to work somewhere the network doesn't reach — a set visit, a flight, a site with no usable uplink. Swarmfile handles that, and there are two different things you can set up before you go:

- **Reading offline** — download a scope so every file in it opens with no network.
- **Writing offline** — *reserve* a scope so it stays editable with no network, and nobody else can change those files while you're gone. This is **Pack & Go**.

You can also just disconnect and keep working — saves queue durably and sync when you're back (see [Casual offline](#) below). Pack & Go is for when you want to be sure the files you need are readable *and* writable while you're away.

The organization's offline access window

One organization setting can bound all offline work: an **offline access window** (owners and admins set it under **Settings → Organization**). It defaults to **7 days** (set it to 0 to turn it off, or choose another window from 1 hour to 30 days). When a window is in force, a computer that hasn't reached Swarmfile within that window **pauses local file access** — reads and writes are refused, even for files kept offline — until it reconnects; the Desktop App explains the pause on the drive, and it clears by itself as soon as the hub is reachable again (no restart needed). It exists so an organization can put a ceiling on how long a machine that stops reaching Swarmfile can keep serving its cached copies. A Pack & Go reservation doesn't override it mid-trip, so make sure the window covers how long you'll be away (lengthen it, up to 30 days, if your team works offline for longer stretches).

Casual offline (no preparation)

If you only need the files you've already been working on, you don't have to do anything special.

- **Saves keep working.** Every write returns immediately against a durable local queue, so a dropped connection doesn't block your application. The queue survives quitting or restarting Swarmfile and resumes when the network is back — see [Reading the sync status](#).
- **Renames, moves and deletes keep working.** A rename or move of a synced file or folder, and deleting a file or empty folder, applies on the drive immediately and is sent to the hub when it's reachable — no waiting for a connection. If the hub later refuses it (a protected branch, say), Swarmfile puts the original name back and reports why. That includes an app that saves by writing a temporary file and renaming it over the original, as most editors do: if that replace is refused, the original file comes back with your unsynced edits to it kept and shown on your drive straight away, and they still upload.
- **Reads work for anything still in your cache.** Files you've opened recently stay local until the cache trims itself, and reading a file you already have open keeps working through a brief outage.
- **Reads do not work for files you never downloaded.** A file you haven't opened yet streams from the hub or a LAN peer, so with no network it isn't there.

The catch is collisions: if a teammate changes a file while you're editing it offline, Swarmfile catches it when your save lands and records a **conflict** rather than silently keeping both. Writes to that file are refused until you resolve it. Pack & Go avoids the surprise by telling everyone else the file is yours before you leave.

Download a scope for reading

Use **Make available offline** when you want the content resident but aren't going to edit it.

- **One file or folder:** in the Desktop App, open the row's menu and choose **Make available offline**.
- **The whole branch:** in the Desktop App, open the **⋮ Project actions** menu on the active project in the left rail and choose **Download entire branch (keep offline)**.
- **From a terminal:** `swarmfile hydrate start [path]`, then `swarmfile hydrate status` to watch it.

While it downloads, the Desktop App shows progress and lets you cancel. When it finishes, the files are pinned against the cache's normal cleanup, so they stay put until you release them.

To undo it:

- **Stop keeping it offline:** **⋮ Project actions** (left rail) → **Stop downloading / keeping the project offline**, or `swarmfile hydrate release [path]`. This drops the offline mark; the content stays in the cache until the cache needs the space.
- **Give the disk space back now:** **Free up space** on a file/folder row, **Free up space (whole project)**... under **⋮ Project actions** in the left rail, or `swarmfile hydrate free-space [path]`. Unsynced changes are never touched — see [Freeing up space](#).

Downloading for reading does **not** stop anyone else from writing to those files, and it does not reserve them for you. That's what Pack & Go adds.

Before you go: reserve a scope (Pack & Go)

Pack & Go does two things at once: it downloads every file in the scope, and it reserves the whole scope with one long-lived offline reservation. The reservation is enforced by the hub, so while it's active nobody else can write those files — and you can keep writing them with no hub at all. The hub refuses a reservation that overlaps one another machine already holds, or that covers a file someone else is editing right now, and names who is in the way. Active reservations are visible in the Desktop App's **Locks** panel, with who holds them and when they expire.

In the Desktop App: open **⋮ Project actions** on the active project in the left rail and choose **Reserve entire branch for offline (Pack & Go)**.... It shows the size before you confirm, lets you pick a reservation length (7, 14, or 30 days), and reports progress; if the scope is blocked, it tells you who holds it and reserves nothing.

From a terminal:

```
# Reserve a folder for a two-week trip: the lease is extended to cover it
swarmfile offline prepare ./Projects/shots/seqA --offline-days 14

# Check progress, resume one a restart interrupted, or release everything
swarmfile offline status
swarmfile offline resume
swarmfile offline return ./Projects/shots/seqA

# Stop a prepare that's still running (reservations already taken are kept)
swarmfile offline cancel
```

`offline prepare` follows the download job to completion; add `--no-wait` to print its id and return immediately (follow it with `offline status`). `offline cancel` stops a running prepare — the reservations already taken are kept, so re-running `prepare` or `resume` picks up from where it stopped.

Reserve only what you need rather than the whole project where you can — the locks you take are locks your teammates can't use, and a smaller scope means a faster prepare.

How long a reservation lasts

- The lease defaults to **7 days** and can't exceed **30 days**.
- It is **only renewed while you're online** — at startup and every six hours after that. A reservation does not extend itself while you're disconnected, so pick a length that covers the whole trip.
- If a lease lapses while you're still offline, the file is no longer exclusively yours: the engine warns you as the expiry approaches, and anyone else can claim the file once it does. Passing `--offline-days` (how long you expect to be away) **extends the lease to cover the trip** when it would otherwise lapse; `--duration-secs` sets an exact lease (up to 30 days) and, if you give an explicit one shorter than the trip, the command warns rather than overriding your choice.
- A single prepare covers up to 100,000 files. A larger scope is refused with a message telling you to narrow it.
- If a prepare is interrupted (a crash, a restart), `swarmfile offline status` names the interrupted scope and `swarmfile offline resume` completes it without re-reserving what you already hold. `offline status` also prints what this machine holds right now (file checkouts and folder/project reservations).
- **Switching a mount to another branch releases the reservation(s) that mount held on the branch you're leaving.** A reservation protects one branch's paths, and after the switch the engine can neither renew nor release it, so Swarmfile hands it back instead of leaving your teammates blocked until it expires. If you'll still need offline work on the other branch, reserve there after switching. **Switching the mount to another project or organization does the same** for the reservations on the project you're leaving. Per-file checkouts are different: they stay yours, keep renewing on the project they were taken in (after a switch within the same organization), and survive quitting and restarting Swarmfile; after a switch to another organization they wait, and resume when you switch back.

Coming back

Choose **Return from offline (release reservations)** under **Project actions** in the left rail, or run `swarmfile offline return` with no path to release every reservation this machine holds. A path releases the reservations at or under that folder. Releasing doesn't delete any content — the files stay in your cache until it needs the space, and you can free them deliberately as above. If saves are still syncing, the release waits a few seconds for them and then tells you how many are still going; they land on their own, but until they do a teammate could take one of those files and force a conflict.

While you're offline

Inside a reserved or downloaded scope:

- Opening and reading files works from the local copy.
- Saving works; saves queue locally and sync when you reconnect.
- Creating files and folders works; the durable create queue lands them when the hub is reachable again.
- Changelists (staged workgroups) keep working.

Anything that needs the hub doesn't: listing someone else's live presence, taking a new lock on a file, comparing history, comment threads, and merge requests all wait until you're back.

Renaming, moving or deleting an already-synced file is the exception: it applies locally and syncs when you're back (see [Casual offline](#)).

Any existing file on a mount you can write to can be edited offline, whether or not it's inside a reserved scope: the save is queued and reconciled when you're back. What a reservation buys you is the guarantee — nobody else can change those files while you're away, so your saves land cleanly. Outside a reservation you may get a **conflict** if a teammate changed the file first (see below); to be sure the files you need are covered, include their folder in the prepare.

Coming back and resolving conflicts

When the network returns, queued saves land on their own. The Desktop App's header stays on **Syncing N** until the queue is empty — see [Reading the sync status](#).

If someone changed a file you edited while you were away, you'll get a **conflict** instead of a silent overwrite. The file is read-only until it's resolved (a save fails with "permission denied", or "access denied" on Windows), and the Desktop App shows it in **Branches & conflicts** in the left rail. From the CLI:

```
swarmfile conflicts          # list blocked files
swarmfile conflicts resolve <path> --auto      # three-way merge when possible
swarmfile conflicts resolve <path> --winner local # or remote, or keep-both
```

`--auto` exits with code 2 when it can't merge cleanly, so a script can tell "a person must choose" apart from a real failure. See [Branches & Merging](#) for the conflict model, and the [CLI reference](#) for every `conflicts` flag.

For administrators

Bulk Pack & Go can be turned off on a machine or fleet with the `SWARMFILE_PACK_AND_GO_POLICY` environment variable on the engine process (`disabled` refuses both bulk prepare and single-file checkout locks). Any value outside `allowed` / `disabled` **fails closed** — a typo disables Pack & Go rather than allowing it — and the engine says so at startup (and `swarmfile-doctor` reports it), so a typo surfaces immediately instead of when someone first tries to pack. On a managed deployment the service can additionally deliver an organization-wide Pack & Go policy to each engine; the stricter of the organization policy and the machine's environment variable applies. See [Engine Config File](#).

Related

- [Working with Files](#) — saving, sync status, freeing up space, and locks.
- [CLI: swarmfile](#) — `hydrate`, `offline`, `fetch`, and `conflicts` in full.
- [Engine Config File](#) — `SWARMFILE_PACK_AND_GO_POLICY` and the streaming/cache settings.

Multiple Mounts

A Swarmfile engine can host more than one mount at a time, so you can work in two projects — or two branches of one project — side by side on the same machine, each as its own drive. This is the quick, human-facing version of the setup described for automated fleets in [Coding Agents](#).

When to open a second mount

- **Two projects at once** — a live project and a prep project, or two clients you switch between all day.
- **Two branches of one project** — compare a look-dev branch against main without switching the whole drive back and forth.
- **An agent or script that needs its own workspace** — a second mount keeps its changelists and locks separate from yours.

Each mount is a full, symmetric mount: its own branch, its own open changelists, its own lock state. Nothing you do in one changes the branch the other is on.

Opening one

In the Desktop App, hover a project in the left rail and click its **+** to open that project as another drive, or use the **+** half of the mount switcher (the drive name at the top of the window), or **+ Open another drive...** inside the switcher. Each opens the **Open a drive** dialog, where you pick:

- the **project** (preselected when you came from a project's **+**) and optionally a **branch**;
- an optional **name** so you can tell drives apart ("agent-2", "client-b") — the same thing the CLI calls a label (`--label`).

Swarmfile chooses where the drive appears — the next free drive letter on Windows, and a new folder beside your main drive on macOS and Linux, named after the drive (or its branch) — and the dialog shows the exact letter or folder before you open it. To pick a specific drive letter or an empty folder of your own, open **Advanced**. If a drive can't open, the dialog says why straight away — for a taken drive letter it offers the next free one — and the failed drive never becomes the current one. A drive that is slow to attach closes the dialog and shows **Connecting...** in the list (it can't be switched to until it's ready); Swarmfile switches to it once it's live, unless you've moved to another drive in the meantime, and a notice says how it went.

From the CLI, the same thing:

```
swarmfile mounts open Y: --label client-b           # Windows: a drive letter
swarmfile mounts open auto --label client-b         # Windows: next free letter
swarmfile mounts open ~/swarmfile/client-b --label client-b # macOS/Linux: a folder
swarmfile mounts open Y: --label dev --branch look-dev-warm # pin this mount to a branch
swarmfile mounts open Y: --label side --project proj_...   # a different project (same org)
swarmfile mounts open ~/agent-2 --exclusive         # exclusive-write mode (agent fleets)
```

On Windows the folder or drive letter must be free — an existing non-empty folder is refused rather than mounted over, while an existing **empty** folder is accepted (the app's folder picker only returns folders that already exist, so this is the normal second-mount shape). On macOS and Linux the mount point is a normal folder, and the same path can't be used by two mounts at once. See [Filesystem Compatibility](#) for the per-platform details.

Keeping two agents off the same file

Locks protect a file against *other people and machines*, but two mounts of one engine share one machine identity — so by default nothing stops two agent mounts from editing the same file at the same time; the later save comes back as a conflict to resolve rather than silently winning. If you want the second writer refused at open instead (and the file protected from rename, delete or atomic-save replace while it's open), open each agent's mount with `--exclusive`:

```
swarmfile mounts open ~/agent-1 --exclusive --label agent-1
swarmfile mounts open ~/agent-2 --exclusive --label agent-2
```

Exclusive-write mode is opt-in per mount. While one exclusive mount has a file open for writing, another exclusive mount on the same project+branch has its write-open refused — with **EAGAIN** on macOS and Linux, and on Windows with the sharing violation apps report as "file in use" — and a rename, delete, or atomic-save replace of that file is refused too, so the holder's file can't be moved, replaced, or removed out from under it. The refusal is recorded (with the holding mount) in the tray's blocked-writes view and in `swarmfile status`. Closing the file — or finishing its save — releases the claim. A mount opened **without** `--exclusive` neither takes nor honors claims, so your human drive and non-exclusive mounts behave exactly as before; the protection is between the mounts that opted in.

In the desktop app the same option is the **Prevent another drive from editing a file this one has open** checkbox under **Advanced** in the **Open a drive** dialog, and an exclusive drive carries a lock badge in the drive list. Like the flag, it is fixed when the drive opens — close and reopen the drive to change it.

Working with several mounts

- `swarmfile mounts list` — every mount, its label, branch, project, and which one is `current`.
- `swarmfile mounts current <id>` — choose which mount commands without an explicit target should act on.
- `swarmfile --mount <id> status` — run any command against one specific mount, without changing the current one. Write the flag **before** the subcommand.
- `swarmfile mounts branch <id> <branch>` — move one mount to another branch without touching the others.
- `swarmfile mounts close <id>` — shut one down. It refuses while a changelist is open (`--force` abandons it), and it always refuses the boot mount.

The Desktop App does the same: each project in the left rail lists its open drives underneath, each row showing the drive's path, branch and status with a control to make it current and a **x** to close it (the boot drive's **x** has no close action — it explains that the drive can't be closed and that switching its project reopens it).

Sparse mounts

Sometimes a mount should show only part of the project — the one folder a coding agent works in, or a single discipline's directory on a giant media project.

`swarmfile mounts open <path> --sparse "<glob>"` opens a **sparse mount** that exposes only the paths matching the include globs; `--exclude` takes the complementary form. The glob dialect matches `materialize --sparse`: a glob matching a directory includes its subtree, a leading `!` excludes an earlier match, and a bare name matches at any depth.

A sparse mount is honest about the boundary: an excluded path is invisible at that drive (no listing entry, an open fails as if it weren't there, and creating or renaming onto it is refused). The scope is per mount, so a second, full mount of the same project on the same machine is unaffected. Offline sets respect the scope —

`swarmfile hydrate start` on a sparse mount can narrow inside it but never widen past it, and Pack & Go takes exactly what the mount exposes. When a render node needs the whole project resident instead, `swarmfile hydrate start` with no path pins everything under the mount's scope.

See the [CLI reference](#) for the full flag set.

Limits and practical notes

- **macOS**: a Mac supports up to 18 simultaneous mounts. If an open fails at that limit, close a mount and try again.
- **One engine, one cache directory**. Don't point two engines at the same cache directory or mount point. Running a separate engine per mount is a valid isolation choice, but it must have its own `SWARMFILE_CACHE_DIR`.
- **Shared scratch cache**. All mounts of the same project on a machine can share one scratch directory for tool caches, so a second mount doesn't re-download dependencies — see [Coding Agents](#).
- **Switching the whole engine** is different from picking a current mount: `swarmfile workspace switch` moves the engine's active org/project — usually in place, without a restart. If the hot path cannot complete, the engine drains and restarts instead (the OS supervisor relaunches it), so the drives briefly unmount and return. Mounts of the project being left follow the switch to the new project's active branch; mounts on other projects stay put — and one on a different project blocks an org switch until it is closed (the app names both before you confirm). See [Switching the active organization or project](#).

Where to go next

- [CLI: swarmfile → mounts](#) — every `mounts` subcommand and flag.
- [Working Offline \(Pack & Go\)](#) — reserving a scope on one mount for offline work.
- [Coding Agents](#) — the headless, one-mount-per-agent version of this setup.

File Previews

Browsing a project shouldn't mean opening a native app just to check you're looking at the right file. The web dashboard renders previews and thumbnails for common media directly, without downloading the whole file to your machine — the preview is generated server-side and you're sent a small image, not the 4 GB source.

This is a dashboard and share-link feature. It doesn't change anything on the mounted drive: there, a file is just a file, and your own applications open it.

Thumbnails in the file browser

The Files view renders a thumbnail for any file type Swarmfile can rasterize, so a folder of renders or plates reads as a contact sheet rather than a list of names. Thumbnails are generated on the server and cached, so the first view of a file pays for generation and every view after is instant.

Thumbnails are produced for:

- **Images** — PNG, JPEG, GIF, WebP, BMP, TIFF, AVIF, and **Photoshop** (`.psd`) files. The PSD path flattens the composite, so you get the image a designer sees, not a layer dump.
- **PDFs** — the first embedded JPEG image found in the file, not a true page render. This works well for the common case of a scanned or photo-heavy PDF, but a PDF with no embedded JPEG (pure vector graphics, text-only, or an embedded image in a different format) gets no thumbnail. Thumbnails use an embedded image; full first-page rendering isn't offered. (The full inline preview below is unaffected — it embeds the real PDF via the browser's own viewer, not this thumbnail path.)
- **Video** — a poster frame (see below).
- **Point clouds** — a top-down preview (see below).

Inline preview

Clicking a file opens an inline preview modal, without leaving the dashboard. Images, PDFs, and plain-text files render directly in the browser. Other types (a PSD, a video, a large TIFF) show their generated thumbnail instead of rendering the raw file inline — the browser can't display the source, but the poster tells you what it is.

Rendering a source file inline is capped at 5 MiB — a browser tab isn't a useful place to view anything larger, and assembling a bigger file just to display it inline holds the server working for no real benefit. Past that size the modal falls back to the type's generated thumbnail or proxy instead (video proxies cover sources up to 16 MiB — see below); for a type with neither, it tells you to download the file. Downloading itself has no size limit.

Anonymous visitors on a [public release page](#) get a smaller inline cap of **256 KiB**, and only images and text render inline there — video and audio stream by range in the player instead (see [Public projects](#)).

Video posters and scrubbable proxies

Video gets two levels of preview, both generated on the server rather than shipping the source anywhere:

- **A poster frame** — a single still, shown as the file's thumbnail, for `.mp4`, `.mov`, `.m4v`, `.webm`, `.mkv`, `.avi`, `.wmv`, `.mpg`, and `.mpeg`.
- **A scrubbable proxy** — a 720p H.264 MP4 the browser can play and scrub through, transcoded in the background from the full source and served with HTTP range requests. This lets a reviewer scrub a cut in the dashboard without pulling the master.

A poster frame is extracted from the first 16 MiB of the source (`SWARMFILE_VIDEO0_POSTER_HEAD_BYTES`), so it works even for a very large video — as long as the frame index sits near the start, which is the usual case for web-friendly MP4s. The scrubbable proxy needs the whole file, so it's generated only for sources up to 16 MiB; a larger video gets no proxy.

Container previews are rate-limited, per project and again across the whole org (500 per project and 2,000 across the org per day by default). An organization's owner can raise or remove the caps under **Settings → Usage → Preview generation budgets** (`0` = no limit); a blank field keeps the default. Once a budget is used up, new video/point-cloud previews are refused with a typed `429` until the next day — the response says whether the project's or the organization's daily budget ran out and when it resets — but already-generated previews are cached and stay free to view. The dashboard flags the refusal with a "Previews are paused" banner rather than leaving a folder of empty thumbnails, and for an owner the banner links straight to those budget controls.

Raw and camera-original formats (`.r3d`, `.braw`, `.mxf`, and similar) are **deliberately excluded** — they aren't reliably decodable without the right vendor codec, so rather than render a wrong or broken frame, Swarmfile shows the generic file icon and leaves them to a native app.

Point-cloud previews

`.las` and `.laz` point clouds get a rendered top-down preview, so a folder of LiDAR or photogrammetry tiles is browsable by eye instead of by filename. Like video, this is generated and cached server-side from the source.

`.e57` and `.ply` are **not previewed** in this version — `.e57` needs a plugin and `.ply` lacks the georeferencing the top-down render depends on. The source cap for a point-cloud preview is 32 MiB.

Previews in share links

A share link carries the same inline preview for images, PDFs, and text, so an external reviewer with no account can see the shared item in their browser. See [Sharing & Collaboration](#) for how share links work, and [Working with Files](#) for the mount side.

What isn't previewed

Limits, so you don't go looking for something that isn't there:

- **End-to-end-encrypted content has no server-side preview.** Previews are generated on the server, and for an E2E project the server cannot decrypt the file — so E2E files show a generic icon, not a thumbnail or proxy. This is the deliberate trade of the E2E tier; see [Security](#).
- **CAD/BIM files (DWG, RVT) show a generic icon; native design-file rendering isn't offered.** Image, PDF, PSD, video, and point-cloud previews are live.
- **Caching is content-addressed.** Editing a file changes its content id, so its preview regenerates on the next view rather than serving a stale one — the first view after an edit pays for generation again.

Search

Search runs on a full-text index and covers entry names (files and folders), project names, and people (principal names). It's a metadata search, not a content search — it finds things by name, not by what's inside a file. Two scoped exceptions: the public [Explore](#) page also matches project descriptions, topics, paths and README text, and the dashboard omnibox searches RFI and submittal title and body text with a project selected (see below).

Where to search from

Dashboard omnibox. The global search bar in the web dashboard is the fastest way to jump to a file, project, or person by name. With a project selected it also searches that project's RFIs and submittals by title and body — see [RFIs & Submittals](#) — filtered to the ones you actually have read access to, the same as everything else it returns; on org-level views (no project selected) it returns name matches only.

Desktop App. The file picker's search box searches the current project's files and folders by name as you type — the quickest way to jump to a file while you're already working in the app. If the hub can't be reached, it falls back to the local cache and labels the results *"from this machine's cache; hub search is unavailable"* so a thin result is explained rather than mysterious.

swarmfile-search CLI. A standalone binary for searching from a terminal or a script. Give it plain text and it searches across entries, projects, and people — an org-wide query (no `--project`) returns projects and people only; file and folder hits need a project scope. Cap results with `--limit`, or force it to search against the local cache or the hub with `--local-only` / `--remote`. Full flag reference: [swarmfile-search](#). The standalone binary stays name-only — to search RFIs and submittals from a terminal, use `swarmfile rfi search` against a mounted project.

Public projects on Explore

Search on the public [Explore](#) page works a little differently: it matches public projects by name, description, topics, org, slug, file paths, and README text, so visitors can find a project by something they remember from inside it. It only ever returns projects that are deliberately published as public.

How it resolves results

The **swarmfile-search CLI** resolves against its local metadata cache first and falls back to the hub, so it keeps working with a flaky or dropped connection to the extent your cache is already populated — a project you've browsed recently stays searchable offline, while one you've never opened won't show up until the hub is reachable again.

Deferred renames show under their projected name. If a rename or move hasn't been confirmed by the hub yet — yours, or one your engine received from a teammate — the local-cache path searches it under the new name: an ancestor rename reprojects full paths, and a match on the old confirmed name is shown under the projected one. The dashboard omnibox searches the hub directly, so it sees confirmed names only. See [Working with Files](#).

The **dashboard omnibox** always searches the hub directly: it runs in your browser, which has no local cache to fall back on, so a slow result means the request is still in flight rather than the index being broken. The panel says **"Searching the hub..."** while a query runs and labels the results as coming from the hub.

Notifications & Inbox

Swarmfile tells you when something needs your attention — an upload failed, a file you have was changed by someone else, someone mentioned you in a comment — rather than leaving you to notice it yourself. Notifications are per-user: everyone on the org gets their own, based on their own activity and preferences.

There are three surfaces, and they show the same events:

- **The web dashboard's notification bell**, with an inbox you can scroll back through.
- **Desktop App toasts** on the desktop, for the same events as they happen on the machine you're working on.
- **swarmfile notifications** from the CLI — `list / unread-count / read / read-all` against the same inbox, for a script or a headless machine that needs to check (or clear) what's waiting without a browser or the Desktop App open. See [the CLI reference](#).

What generates a notification

You'll get a notification for:

- **A failed upload** — a save that couldn't be uploaded after retries. Because saving is asynchronous, this is how you find out something didn't land from anywhere; in the Desktop App the header also counts failed changes until they're dealt with (see [Working with Files](#)).
- **Quarantine** — the *applied* notice goes to the org's owners, deliberately not to the quarantined account; the person released gets the release notice. If your own saves start being refused, the Desktop App's "Why did my save fail?" shows the reason. See [Security](#) for what triggers quarantine.
- **A quota or billing block on upload** — an upload refused because the org hit a storage limit or has a lapsed subscription, so you know why writes started failing rather than seeing an opaque error.
- **A file you have was changed by someone else** — the project-wide "someone edited this" signal.
- **Comments and @-mentions** — a new comment on a file, or someone tagging you specifically.
- **Watched files and folders** — a change to something you explicitly subscribed to with [watch](#).
- **Unlock requests** — someone asking you to release a lock on a file you're holding.
- **Merge Request activity** — one opened, reviewed (approved, changes requested, or just commented on), merged, closed, or reopened; a draft becoming ready for review; you being assigned, asked to review, or unassigned (and a review request withdrawn); and labels being added or removed. See [Branches & Merging](#).
- **A hosted-CI run failing** — the run's starter is told when a run ends in failure (a timeout, a cancel, a queue expiry, a blocked start or an infrastructure fault included); on a protected branch the branch's recent committers get it too. Bell/tray and the personal webhook, not email. See [Hosted CI](#).
- **A project nearing its storage limit** — sent to the org's owners and admins when a project's file information and history pass the early-warning mark, again if new files and versions get paused at the limit, and once more when it recovers. The project's **Metadata** tab explains what to do. See [Project storage limit](#).
- **A teammate asking for git access** — sent to the org's owners and admins when someone tries `git clone / git push` on a project that hasn't been enabled for git yet. It arrives once per project, not once per attempt, and names the fix: [Clone a project with git](#).
- **Your org's dedicated storage finishing provisioning** — sent to the person who enabled the switch (or created the first project on a paid plan) when the bucket is live, since that can complete well after the request returned. See [Dedicated Storage Isolation](#).
- **Your org's request budgets or spend cap changing** — the org's owners and admins get a bell/tray notice when an owner or admin changes the org's request budgets or [spend cap](#), including when a spend-cap threshold is crossed or a metered write is refused. Bell/tray only; the Activity feed and audit log are the trail.
- **A commit dispute**: two derivations of the same commit disagree. The org's owners and admins are told (bell/tray only) while git holds that commit back, and told again when Swarmfile settles it on its own, usually within a minute or two. When a differing derivation arrives for a commit that is already confirmed, the notice says the confirmed commit is still served and nothing is held back. See [Clone a project with git](#).
- **An org webhook being auto-disabled** — when a webhook is turned off after repeated delivery failures, the org's owners and admins are told by bell and email. See [Webhooks](#).

One event is desktop-only: when a found update has sat uninstalled for a week, the Desktop App raises a one-time notification reminding you. It doesn't appear in the bell or by email, and nothing ever installs on its own — see [Release Channels & Updates](#).

Email notifications

Some notifications can also be delivered by email, so an important one reaches you when you don't have the dashboard open or the Desktop App running. The email-able events are the ones you'd want to know about while away from your desk: **failed uploads, quarantine (the applied notice to owners, the release to the affected user), a quota/billing block, a file changed by someone else, a public access request or public RFI filed on a project you own or administer, comments, @-mentions, being assigned to a merge request, having your review requested on one, having changes requested on one you opened, RFI/submittal activity (created, assigned, response posted, revisions requested, due soon, overdue, closed, voided, and the ball-in-court digest), a project nearing or hitting its storage limit, an org webhook being auto-disabled, and your org's dedicated storage finishing provisioning**. Every other Merge Request event — it being opened, approved, merged, closed, or reopened — stays bell/toast-only; those happen often enough on an active project that emailing every one of them would be noise, not signal.

Watched-file changes and unlock requests are **in-app only** — they show in the bell and as Desktop App toasts, but aren't emailed.

Email is on by default for the events listed above, with two exceptions: "**a file you have was changed by someone else**" and **new comments on a file** ship muted — on an active project those fire for ordinary teammate activity, and emailing every one of them is noise. The bell still records both. You control every kind individually, so you can turn those two on, or mute others; mentions, failed uploads, quota blocks, quarantine and the other attention kinds stay on unless you mute them.

Quiet hours

You can set a quiet-hours window — a start and end time in your own timezone — during which we won't email you, so an overnight render failure doesn't wake you. A notification raised during quiet hours still appears in the in-app inbox immediately; only the email is suppressed for the window.

Where to configure it

Email preferences live in the web dashboard's **Settings → Notifications**: the master on/off switch, quiet hours, your timezone, and the per-kind toggles (including the RFI/submittal kinds, each of which can be muted without touching the rest). The table lists every notification kind — kinds that never send email (watched-file

changes, unlock requests, MR opened/merged/closed/reopened, CI runs, and the other bell-only events listed above) show a dash in the Email column, and a configured personal webhook is their per-kind control. Set them once and they apply to your account across every machine you sign in on. `swarmfile notifications preferences get / set` reads and writes the on/off switch, quiet hours and timezone from the CLI (clearing either of the latter is dashboard-only — the CLI has no null form), and `preferences set --kind file_changed_by_other=off` sets a per-kind override (repeat `--kind` to set several; `default` clears that kind's override — email and webhook — back to the kind's defaults: the master switch plus that kind's own default, which is off for the two high-volume kinds above and on for the rest). The dashboard remains the friendliest place to see them all at once.

Personal webhook

A third delivery surface, alongside the bell and email: a signed HTTP callback to a URL only you configure, with the same per-kind mute controls email already has. It's a fire-and-forget best-effort channel — no delivery log or retry, held to the same bar as email rather than the fuller treatment an org-wide webhook gets. See [Webhooks](#) for setup and how to verify the signature. One asymmetry worth knowing: a per-kind email override turns that kind's email on even when the email master switch is off, but the webhook's own on/off switch is a hard gate — its per-kind rows only apply while the webhook is enabled.

Notifications vs. the activity feed

The notification inbox is **yours** — the events relevant to you personally. It's distinct from the org **Activity** feed, an org-wide audit log open to any member (ACL-filtered — you see events on entries you can read; quarantine incidents and org-policy events stay owner-only within it). Its CSV export is more restricted still: owner-only, and it needs the Pro plan's `audit_log` entitlement. If you're looking for a complete record rather than your own attention queue, that's [Operations → Audit log](#), not this.

That Activity feed is the same stream wherever you meet it — three surfaces onto one feed:

- **The web dashboard's Activity view** — the full, scrollable, ACL-filtered record described above.
- **The Desktop App's activity feed** — the right-rail, cross-machine feed of what's happening in the project you're currently working in. Scoped to the active project rather than the whole org, it's the at-a-glance version for while you're at your desk. Runs of the same person doing the same thing collapse into one row you can expand; updates to system files (`.DS_Store`, Office lock and temp files, macOS safe-save `*.sb-...` leftovers) stay hidden behind a **Show system-file updates** toggle rather than being deleted. If you select a file, the feed narrows to that file's events.
- **swarmfile activity (`list` / `export` / `tail`)** — the terminal equivalent: `list` is point-in-time, `tail` follows the feed live (by polling), and `export` writes the owner-only CSV. See [the CLI reference](#).

Same events underneath; pick whichever surface fits what you're doing.

Not to be confused with `swarmfile tail`, which streams the engine's **local** event bus for the mounted project (comments, watched-file changes, lock conflicts) rather than this org-wide Activity feed.

Where to go next

- [Sharing & Collaboration](#) — comments, mentions, and watch, which are what generate most notifications.
- [Working with Files](#) — asynchronous saves and the upload queue, which the failed-upload notification reports on.
- [Webhooks](#) — the personal webhook channel in full, plus the org-wide version an owner can set up.

Sync Exclusions

A mounted project normally syncs everything: any file you write shows up for every other machine mounting it. That's the right default, but not everything belongs on the hub — `node_modules/`, build output, local caches. Swarmfile reads `.gitignore` the same way git does and treats a match the same way sync does: the file stays exactly where you put it and works normally on your machine, it just never uploads.

This is on by default, the same way a `.gitignore` in a fresh git repo is just expected to work. When you create a project from one of the ready-made templates (Media, Git, Revit, Geospatial), Swarmfile seeds a starter ignore file for that workflow — edit or delete it like any other file.

`.gitignore` VS `.swarmfileignore`

If your project is already a git repo, its `.gitignore` is enough — Swarmfile reads the same file, with the same syntax, and nothing extra to maintain. `.swarmfileignore` exists for two other cases: a project with no git repo at all (a CAD or media project versioned only by Swarmfile itself), and excludes that are only about sync, not about git — a local scratch folder or a cache directory you'd never want in a git commit either, but don't want to mix into a `.gitignore` that other tooling also reads. The two files use identical pattern syntax and combine additively: a directory can have both, and either one can exclude a path the other doesn't.

Nested rules, just like git

Rules cascade the way `git help gitignore` describes: a `.gitignore` in a subdirectory is layered on top of the project root's, and if the two disagree, the deeper one wins — including a `!pattern` negation un-ignoring something a shallower rule excluded. `.git/info/exclude` and your global `core.excludesFile` are honored too, at the bottom of the precedence order: global excludes, then `.git/info/exclude`, then the project root's own `.gitignore` / `.swarmfileignore`, then anything closer to the file. Within one directory, a `.swarmfileignore` line is applied after that directory's `.gitignore` lines, so it can override a rule the `.gitignore` set for the same directory.

What a teammate actually sees

An ignored file isn't deleted or blocked — it behaves exactly like an untracked file in git. You can create it, write to it, and read it back on the machine that made it, same as anything else on the mount. What changes is visibility to everyone else: the file never uploads, so it doesn't appear in `ls`, Explorer, Finder, or the web dashboard's Files view on any machine whose own ignore rules match it. It isn't hidden metadata with empty content — it's absent from the listing entirely, the same as a file that was simply never created there.

There's one exception worth knowing: opening the exact path directly (not browsing to it) still works, but only if the bytes are already sitting on that machine — because this machine wrote them, for instance. If they aren't there, Swarmfile won't fetch them from the hub just because you asked for the path by name; an ignored file is never fetched, only ever served from what's already local.

The web dashboard's Files view has a **Show ignored files** toggle for anyone who needs to see the unfiltered truth — an admin auditing what's actually sitting on the hub, or checking whether something is worth cleaning up (see below). With it off (the default), the dashboard matches whatever an ordinary mount would show you.

Ignore rules and releases

Sync-ignore hides paths from listings and stops future syncs — it does **not** keep anything out of a release. A release publishes the whole tree of its tag, and a path that's already on the hub (including one that synced before your rule existed) still goes in. The only file that keeps paths out of a release is the committed `.swarmfile/publicignore.yml`. The publish preflight calls out exactly this gap before anything is public: tracked paths your `.gitignore` / `.swarmfileignore` hide that would still be published, by name.

If a file was already synced before you ignored it

Ignoring a path doesn't retroactively unsync it. A file that made it to the hub before your `.gitignore` rule existed — or one created by a teammate's machine that doesn't have the matching rule — keeps its existing hub-side content indefinitely. This is exactly how git treats a file you stop tracking with `git rm --cached`: the ignore rule stops *future* changes from syncing, it doesn't erase what's already there.

To actually remove that old content, run:

```
swarmfile ignore-clean
```

It scans your project (or a subtree, if you pass a path) for anything already on the hub that your current ignore rules would exclude, shows you exactly what it found before touching anything:

```
3 item(s) already on the hub match your ignore rules, totaling 340.0 MiB:
  node_modules/ (directory)
  dist/ (directory)
  .env.local (file, 2.1 KiB)
This will remove them from the hub for everyone – recoverable from Trash for a while, not instant or permanent.
Proceed? [y/N]
```

and only deletes after you confirm (or pass `--yes`). Deletion goes through the same Trash system as any other delete, so it's recoverable, not instant or permanent — but it does remove that content for *everyone*, not just you. Be careful running this on a shared project if a teammate's machine might not have picked up the same `.gitignore` yet, or might be relying on something inside a directory you're about to clean up: `ignore-clean` only knows about *your* machine's current ignore rules, not anyone else's. It won't touch a path currently locked by someone else — the delete is refused, not silently skipped — but that only catches a file open for editing right now, not one a teammate's tooling reads without holding a lock.

macOS metadata files never sync

On macOS, two kinds of file the system writes by itself stay on your machine without any ignore rule: AppleDouble `._name` files (where macOS keeps a file's Finder tags, comments, and other extended attributes on the drive) and `.DS_Store` (Finder's per-folder view settings). They're meaningless to anyone else, and syncing them would put one extra hub entry beside nearly every file and folder a Mac user touches. They're hidden from listings on the machine that wrote them and never upload. This applies even with sync exclusion turned off, and no whitelist rule can override it.

Two limits keep it from touching anyone's real data. It only applies on macOS: on Windows or Linux, a `._x` or `.DS_Store` you create came from an unzipped archive or an imported repo, so it syncs like any other file. It also only applies when the file is created: one already on the hub stays listed, readable, and synced.

Office owner files never sync

Word, Excel, and PowerPoint write a small `~$name` owner file beside a document while it's open — it records who has the document open, and the app deletes it when the document closes. Swarmfile keeps those files on the machine that wrote them and never syncs them, with no ignore rule: it applies even with sync exclusion turned off, and no whitelist rule can override it. Unlike the macOS files above, this applies on every platform Office runs on. It's also create-time only, the same as the macOS rule — an owner file that already reached the hub stays listed and synced.

One thing this deliberately does not cover: Word's `~WRD####.tmp` / `~WRL####.tmp` save scratch files still sync. Word renames those over the document as part of saving, so keeping them local-only would strand the save; they normally disappear the moment a save completes. If a save fails or Word crashes, a leftover scratch file can show up as an ordinary project file until you delete it.

macOS safe-save scratch files stay local

When an editor on macOS saves a document, it may stage the save in a scratch folder named `<name>.sb-<8 hex>-<6 letters or digits>` beside the file (TextEdit's save is the usual source: it creates the folder, writes the new copy inside it, renames that copy over the document, then removes the folder). Swarmfile recognises that exact shape and holds its hub registration back long enough for a normal save to remove the folder first, so a routine save never publishes one. This isn't an exclusion rule — it isn't governed by `.gitignore` / `.swarmfileignore`, and no `!` line turns it off.

Like the Office rule above it is create-time only, with the same caveat: a save that stalls past that window, or an app that crashes mid-save, can leave a scratch folder that surfaces as an ordinary project file — empty and unopenable, named like `report.docx.sb-1a2b3c4d-Ab12Cd` — until the app's next save removes it or you delete it. Listings and the Activity feed hide the shape behind **Show system files** either way.

Git metadata never syncs

A repository's `.git` directory is machine-local metadata, and syncing it is pure churn: packfiles, refs and hook samples change on every git operation, each change becomes its own hub entry and upload, and no teammate's machine can use another machine's checkout state. So Swarmfile treats `.git` — the directory, or the `.git` file a linked worktree or submodule uses as its gitlink — as **local-only on every platform**, with no ignore rule and with sync exclusion turned off: `git init`, `git clone` and every commit inside a mounted project stay on your machine, and nothing a current build puts under `.git` is ever uploaded, staged for publishing, or listed to teammates. No `!` rule can unmask it. Only the exact name `.git` is matched — `.gitignore`, `.gitattributes`, `.gitmodules` and `.github` are ordinary project files and sync like any other.

One move caveat: a file that is already synced can't be moved *into* `.git`, and a synced entry can't be renamed to the name `.git` — the operation is refused (an invalid-argument error), because there is no hub `.git` to move it under and a synced entry can't become local-only. Create or save the file inside `.git` and it stays local, as above.

`.git` content that already reached the hub keeps its stored copy until it is cleaned up, the same as any file that synced before an ignore rule existed. While sync exclusion is on (the default) it is no longer shown to mounts or the web dashboard and stops syncing; with the switch off, pre-existing content behaves exactly as the other built-ins' does — it stays listed and keeps syncing. Turn sync exclusion back on to hide it and clean it: `swarmfile ignore-clean` lists those `.git` trees as candidates alongside anything your own ignore rules match, so a one-time cleanup removes them for everyone. `swarmfile check-ignore .git/HEAD` reports it as `ignored – excluded by built-in rule (git metadata): '.git'`.

Windows Mark of the Web never syncs

Windows tags a file you download with a `Zone.Identifier` alternate data stream, the "Mark of the Web": it records that this machine got the file from the internet, which is what makes Office open it in Protected View and Windows warn before running it. That is about how *your* machine got the file, not about the file, so Swarmfile keeps it on the machine that wrote it and never syncs it. Copying a downloaded file onto the drive works as usual, and the stream reads back on that machine, but your mark is never sent to a teammate. Other alternate data streams still sync. Like the rules above, this needs no ignore rule and nothing overrides it. Unlike them, it also covers marks already on the hub: a stored mark is never copied onto a machine that doesn't already have it, and the hub deletes those stored copies on its own. One a machine already picked up from the hub stays there, though, and just stops syncing: Swarmfile can't tell it apart from that machine's own mark, and removing it would remove a security warning, so it errs toward keeping the warning.

Rules apply across the team, with one caveat

Because `.gitignore` and `.swarmfileignore` are themselves ordinary synced files, everyone mounting the project ends up with the same exclusion rules once those files have synced to their machine — you don't need to hand-distribute them. The one moment this isn't instant: on a machine that just mounted the project and hasn't fetched anything yet, there's a brief window where a not-yet-fetched ignore file contributes no rules, so sync proceeds normally until it arrives in the background. This resolves itself within moments, not something to work around.

Checking whether a path is ignored

If a file isn't showing up where you expect — or one you meant to exclude keeps syncing anyway — ask directly instead of guessing:

```
swarmfile check-ignore ./Projects/node_modules/pkg.js
```

```
ignored — excluded by node_modules ignore rules: `node_modules/`
```

For a path nothing excludes, it just says `not ignored`. The output names both which directory's rules made the call and the exact pattern that matched — useful when a nested `.gitignore` is overriding one closer to the project root, or a `!negation` is un-ignoring something you expected to stay hidden.

Turning it off

Sync exclusion is on by default. Turn it off from the Desktop App's Settings panel (**Sync-exclude .gitignore'd files**), or set `sync_ignore_enabled` to `false` in that machine's `config.json`, or export `SWARMFILE_SYNC_IGNORE_ENABLED=0`. With it off, no path is ever hidden from listings, excluded from upload, or refused on read by a `.gitignore` or `.swarmfileignore` rule. The built-ins are not governed by this switch where they act: a newly created macOS metadata file, Office owner file or `.git` entry is local-only whichever way it points, and Windows' Mark of the Web is never synced — but content already on the hub follows each section's own note above.

Flipping this switch — either direction — restarts the engine. It changes how the mount itself behaves, not something that can be swapped out from under it while it's running.

Example

A typical `.gitignore` at a project's root:

```
node_modules/  
dist/  
*.log
```

With sync exclusion on, `node_modules/` and `dist/` build up locally as normal — `npm install` and your build both work exactly as they would on disk — but nothing under them ever uploads, and teammates never see those directories appear in their own mount. (`.git/` needs no line here: git metadata never syncs, on any platform — see above.)

Performance & Disk Usage

Swarmfile is fast by default for the two things people notice — opening files they have touched before, and saving without waiting. When it isn't, it's almost always one of two resources: local disk (the cache) or the network path a file took (LAN peer, seed node, or cloud). This page covers both, plus how to keep the cache from growing into a laptop's free space.

The local cache

Every file you read is copied into a local block cache and kept there until space is needed. The cache is **not** a copy of the project — it holds what you have actually touched, newest first, and trims itself as it fills.

- **Default size: 10 GiB**, per machine and per user.
- Change it in the Desktop App's settings, or from a terminal: `swarmfile cache set 50` (GiB), or `swarmfile cache status`.
- The mount always reports **at least 1 GiB free** to applications, so an app's "disk full" is about your actual local disk, not the cache being full-but-evictable.
- A **seed/NAS node** is the exception: it fetches whole projects (mirroring them to cloud storage) and keeps them in a warm cache with its own budget. Size `SWARMFILE_CACHE_MAX_BYTES` for the working set you want resident — the sample seed config suggests 100 GiB as a starting point — and remember that a block evicted past the budget simply re-fetches on the next read.

Files you want resident on purpose — a render source tree, a site archive — should be pinned rather than left to the LRU:

```
swarmfile hydrate start ./shots/seq010 --yes # download and pin (no prompt)
swarmfile hydrate status                    # progress of the current job
swarmfile hydrate release ./shots/seq010    # unpin when done
```

Pinned content can exceed the cache cap; it is never evicted, so don't pin more than the disk can hold.

Making reads fast

Reads come from, in order of preference: the local cache, a peer that has the file (**LAN first**; then **same-office seed nodes and peers**, then **other-office seeds and peers**, each bucket sorted by measured round-trip time), then cloud storage (via a presigned URL direct to the org's bucket — not through the hub). Two things follow:

- **The first person to open a file pays for it; everyone else in the office gets it over the LAN.** If transfers look like they are coming from the cloud when peers exist, check `swarmfile status` for the peer count, then run `swarmfile-doctor`: the `mdns listen` and `host firewall` checks name the usual causes.
- **Erasur coding spreads the load where it applies.** Saved blocks large enough to be sharded (≥ 64 KiB) on WAN-heavy or cloud-backed saves are Reed-Solomon 10+4, so a read of those can reconstruct from any 10 of 14 shards instead of stalling on one machine, and a slow or offline peer costs redundancy rather than the read. The office LAN path usually fetches the stored chunk directly — erasure coding is skipped when two or more LAN peers already hold it — so this applies to the saves that used it. With [deliberate shard placement](#) on a seed-node roster, those shards land on guaranteed homes across the office instead of wherever caching happened to put them. The exact rules are in [Storage Format](#).
- **Pre-hydrate the hot set.** For a render farm or a shoot day, `hydrate start` before the job starts beats streaming during it. `swarmfile uploads --wait` is the matching pre-flight when you need an upload to have landed before reading it elsewhere.

For a single very large file, ask for the part you need instead of the whole thing:

```
swarmfile fetch ./plate.exr --tail 209715200 # last 200 MB
swarmfile fetch ./plate.exr --follow          # track a moving reader (playback)
```

The engine also **reads ahead** as you go. Video containers (`.mp4` , `.mov` , camera and NLE media) start on the first sequential read and keep roughly 15 seconds of playback buffered ahead, up to 64 MiB; other file types ramp from the third sequential read, once about 2 MiB has been covered, with a 16 MiB ceiling. Scrubbing is bounded too: a seek drops the window to one small batch and it only widens again after steady forward reading, so dragging a timeline does not pull a full window per stop. A cold, cloud-only first open spends a second or two building that lead before it settles; a warm cache or a LAN peer removes the wait. Tune it in the Desktop App's **Settings → Read-ahead** (seconds; blank = the file type's default — a larger lead keeps high-bitrate playback smoother at the cost of fetching more ahead).

Metadata-heavy workloads (thousands of small files, or an app that stats before it opens) benefit from a larger attribute cache: `swarmfile attr-cache set 10` (seconds; macOS/Linux — the engine restarts to apply it, and Windows mounts are unaffected). CAD tools that open external references can leave xref prefetch on (the default) or turn it off for a pathological reference graph — see [Environment Variable Index](#).

Where the bytes came from

To see that preference order in action — on a live project, per file — open the Desktop App's **Data sources** panel (left rail → **Data sources**). It shows, for each file and in total, how much of what you read came from **this computer**, from **a colleague's computer** (a LAN or WAN peer), and from **the cloud**, with first-read times. A read-ahead line underneath reports how much the engine fetched speculatively, how much of that was actually read, how many seeks (run resets) the window saw, and the file with the most unread bytes — the same numbers `swarmfile doctor` prints as **read-ahead efficiency** (doctor additionally reports the seek-storm share and its depth, which the panel omits); a window that seeked without prefetching shows **Read-ahead idle** beside its seek count. **Reset counters** starts a fresh measurement window; the process-wide bandwidth numbers restart with it. A one-line version sits in the header's **Status** popover.

The same numbers are on the command line — the form to use when they need to end up in a spreadsheet or a bug report:

```

swarmfile stats          # per-file table
swarmfile stats --json   # full report: version, machine, mount path/branch, sizes, timings, bytes by source
swarmfile stats --csv    # one row per file for a spreadsheet
swarmfile stats --reset  # zero the counters after printing

```

For a genuinely cold start — a benchmark, a demo, or "is the LAN actually being used?" — `swarmfile demo reset` empties this machine's cache for one project (the same whole-project sweep as **Free up space**, including offline pins; never unsynced saves) and zeroes those counters, so the next run measures from zero. Run it on every machine involved; see [swarmfile demo reset](#).

Bandwidth and schedules

Throttling is **off by default**. If your office link is shared and Swarmfile's background traffic is crowding out people's calls, turn it on with a real link size first:

```

swarmfile throttle enable \
  --download-mbps 500 --upload-mbps 100 \
  --office-hours-start 08:00 --office-hours-end 18:00 \
  --office-hours-wan-pct 20 --off-hours-wan-pct 80

```

That allows 20% of the link to WAN transfers during office hours and 80% outside them; LAN traffic and app-requested reads are not held back by the schedule. Status and changes: `swarmfile throttle status` / `swarmfile throttle disable`.

The same control lives in the Desktop App: **Settings** → **General** → **Limit bandwidth**, and the tray menu's **Bandwidth...** item opens it directly. (The tray item is disabled when the engine was built without the throttle feature.) The same destination answers the `swarmfile://bandwidth` deep link, so a shortcut or a link you send a colleague lands on the throttle card rather than the top of Settings. It is per computer, not per organization.

Two behaviors worth knowing:

- Turning throttling on with **no bandwidth set** applies a guessed symmetric cap (about 100 Mbps), so always set the link size.
- Interactive traffic (opens, small reads, saves' control calls) is marked for priority over bulk transfer (DSCP, when your network honors it). You rarely need to touch the DSCP values; they exist for networks that prioritize by tag — see the [Environment Variable Index](#).

Uploads are designed not to block

A save returns as soon as the bytes are safely queued locally; the upload continues in the background, survives restarts, and resumes by itself. A 500 GB file on a 100 Mbps uplink will take days to reach the cloud, and that is expected — not a stuck sync. `swarmfile uploads` shows what is still in flight (across all machines, not just yours), and `--wait` blocks until nothing is. Other people can fetch the part they need before the upload completes — see [While a large file is still uploading](#).

If a save isn't landing, it is on the **stuck** list, not the upload list: `swarmfile sync stuck` says why and for how long, and the Desktop App's header says **N changes not syncing** or **N changes failed**. Nothing in the queue is silently dropped: a transient hub refusal backs off and retries, while a refusal a person can clear — plan, membership, auth policy, an archived project, root-create authority — is *held* until the block clears. A removed member's own unsynced saves stay on their machine (ciphertext once the key is dropped) and upload if access returns.

A burst of many small files is a different shape from one large one. Each save needs a handful of control-plane calls, and Swarmfile paces them to the hub's request allowance for your plan, so a mass overwrite of hundreds of files converges over minutes while the header counts up to **Synced**. That is deliberate pacing, not a stall — the work is durable throughout, and a restart resumes it.

Freeing up local disk

In rough order of how much they return:

- Let the cache do it.** Lowering the cap (`swarmfile cache set 5`) applies immediately and evicts unpinned content down to the new size, no restart. (`swarmfile cache reclaim` looks similar but isn't a general eviction command: it removes an old cache directory layout that is otherwise only kept as a read-only fallback, and it has a `--dry-run` to preview.)
- Release deliberate pins** you no longer need: `swarmfile hydrate release <path>` followed by `swarmfile hydrate free-space <path>` (or **Free up space** in the tray). Free up space never touches unsynced changes.
- Remove content that should never have synced.** If a project synced `node_modules/` or render caches before an ignore rule existed, `swarmfile check-ignore <path>` confirms the rule, and `swarmfile ignore-clean` removes the already-synced content (a soft delete; recoverable from Trash).
- Empty Trash and old branches.** `swarmfile trash purge <entry-id>` and `swarmfile branch prune` shrink the **hub's** storage; freeing server-side space is not about your laptop, but it matters for the org's allowance.
- Build caches are separate.** Tool caches live in the project scratch directory, not the block cache: `swarmfile scratch gc` reclaims stale ones, and `swarmfile mounts scratch-dir <id>` prints where they are. See [Coding Agents](#).

If the Desktop App's disk usage is still surprising, check the size of the cache directory itself (`~/ .cache/swarmfile`, `%LOCALAPPDATA%\swarmfile`) and compare it to the configured cap — pinned/hydrated content legitimately exceeds the cap, and `swarmfile hydrate status` lists what is pinned.

Diagnosing "everything is slow"

- `swarmfile status` — peers in the office, in-flight uploads, commits waiting, stuck saves.
- `swarmfile doctor` — DNS, hub reachability, mDNS, firewall, and the LAN-consent checks.
- Check whether the slowness is **first open** (network) or **every open** (local disk or an app doing its own scanning). First opens are expected to be slow for a file nobody nearby has; repeat opens should be cache-speed.

4. Very large projects: a project with millions of files and versions can make full-project browsing and history feel heavier — history retention and project splitting are in [Operations](#).

Where to go next

- [Working with Files](#) — sync status, freeing up space, and watching an upload.
- [Working Offline \(Pack & Go\)](#) — reservations for field work.
- [Deployment Topologies](#) — seed nodes and LAN-first office design.
- [Environment Variable Index](#) — every cache, streaming, and QoS tunable.

Version Control

Swarmfile versions every file the way a VCS versions code — except the files are whole binaries (video masters, CAD assemblies, GIS rasters) instead of text. There's no meaningful line-level diff for a 200GB ProRes file or a Revit model, so the model is built around whole-file commits and whole-file history instead of line hunks. If you're used to git or Perforce, the mental model carries over; the mechanics underneath don't.

Changesets

A **changeset** is a named, revertable commit that can span multiple files, ordered in a DAG the same way git commits are — you can always roll every file in it back to its pre-commit state as one unit. A revert or point-in-time restore (`checkout`) covering more than the hub's interactive limit (about 5,000 entries) runs as a durable background job: the CLI waits for it by default (Ctrl-C or `--no-wait` detaches with exit 3, followed with `swarmfile op status`), and the result appears in the log once it lands. Whether the underlying saves actually land *atomically* (all at once, invisible until they do) depends on which of the two modes below produced it.

Swarmfile supports two ways of producing changesets, and you pick the one that matches how you work.

Mode A: auto-landed (default)

By default, your saves land as commits automatically as you work, and immediately: each save is live on HEAD the moment it uploads, visible to teammates right away — it isn't held back waiting on anything else. What Swarmfile groups is the *changeset* those saves get filed under: consecutive saves are batched into one open changeset while you're actively working, and that changeset closes and commits once saves go idle for `changeset_idle_secs` (default 5 minutes, configurable via `SWARMFILE_CHANGESET_IDLE_SECS` or the config file). A burst of edits becomes one named, revertable commit; the next burst after a break becomes another. You don't open or close anything yourself — this is the lower-friction mode and it's what most people want most of the time. One exception: a save whose bytes are identical to the version already on HEAD files nothing — there's no change to commit, so it doesn't create a new version (an app that rewrites a file unchanged on close therefore can't spam history).

```
swarmfile commit -m "Color grade pass 2 on interview B-roll"
```

`commit` opens and submits a changeset in a single step — the Mode A pattern, useful when you want to attach a specific message to a specific save rather than relying on the automatic one.

Mode B: staged, explicit submit

For work you want to land as one deliberate, reviewed unit — a multi-file asset swap, a batch of renders that only make sense together — open a changelist, stage saves into it, and submit (or cancel) the whole thing as one atomic transaction. It's the same shape as a git index: nothing is visible to teammates until you submit, unlike Mode A where each save is already live the moment it happens.

```
swarmfile changelist open -m "Swap all Act 2 plates to final grade"
# ...work happens, saves are staged into the open changelist...
swarmfile changelist status
swarmfile changelist submit
```

`changelist submit -m "message"` sends the commit message at submit time (git-style); without `-m` the group's open-time label is kept, and a group with neither is committed without a message. Submitting checks write access to every file in the changelist at that moment, so a save staged before a permission change doesn't bypass it; a refused submit leaves the changelist staged, nothing lost.

If you change your mind partway through, `swarmfile changelist cancel` discards the staged changes instead of submitting them.

On the Desktop App, the **Changes** tab shows the same work with a git-style commit box: type a summary (and an optional description) and press **Publish**. The message is required, just like a git commit, and is sent with the publish. A group's open-time label, if it has one, prefills the summary. The switcher's per-row publish (acting on another open group) sends no message — that group keeps its label, or is committed without a message when it has none.

Multiple changelists at once

You're not limited to one open changelist. Open a second while the first is still open, and each save stages into whichever one is *current* — the same idea as Perforce's `-c` flag:

```
swarmfile changelist open -m "Swap all Act 2 plates to final grade"
swarmfile changelist open -m "Quick fix: wrong LUT on shot 12"
# saves now stage into the second changelist, since it's current
swarmfile changelist list
# #a1b2c3d4 "Swap all Act 2 plates to final grade" - 6 file(s) staged
# * #e5f6a7b8 "Quick fix: wrong LUT on shot 12" - 1 file(s) staged
swarmfile changelist current a1b2c3d4
```

`changelist list` shows every open changelist with a `*` marking which one is current; `changelist current <id>` switches. This lets you juggle two unrelated in-flight edits without submitting one to start the next — no need to interrupt a large staged batch just to make an unrelated quick fix.

If a file was staged into the wrong changelist, move it without unstaging and restaging by hand:

```
swarmfile changelist reopen "Renders/Act2/shot12.mov" --to e5f6a7b8
swarmfile changelist reopen --to e5f6a7b8 --from a1b2c3d4
```

The first form moves a single staged file; the second (no path) moves every staged file from one changelist to another in one atomic step. Either can target a specific changelist without switching to it first — `submit` / `cancel` also both accept `--id <changeset>` for the same reason.

Setting work aside

If you need to leave the branch or project but aren't ready to submit or cancel, **park** the work instead — it's set aside and kept alive on the hub, then restored when you come back to that scope or explicitly:

```
swarmfile changelist park      # set every open changelist aside
swarmfile changelist parked    # list parked work, with age and file count
swarmfile changelist resume    # restore parked work for this branch
swarmfile changelist discard   # abandon it (asks first; --yes skips)
```

Parked work is a **set, not a stack**: `resume` restores *all* parked work for the branch, and `discard` abandons all of it (or one, with `--changeset <id>`). A branch or workspace `switch` into a scope with staged work offers "Park for later", which does the same thing before the scope flips; the Desktop App's **Parked changes** panel lists, resumes, and discards the same work.

If you're coming from git, `swarmfile stash` maps here: `stash` → `park`, `stash pop` → `resume`, `stash list` → `parked`, `stash drop` → `discard`.

Viewing history

`swarmfile changesets list` shows the commit log, newest first. `swarmfile log` gives you the same history as a compact colored feed, closer to `git log --oneline`, when you just want a quick scan.

Checkout and restore

`swarmfile restore` restores a scope to how it looked at a past commit — the whole project, a folder, or a single entry:

```
swarmfile restore 482 --path "Renders/Act2/final.mov"
swarmfile restore HEAD~3
```

The commit can be a number (`#N` from `swarmfile log`), `HEAD`, `HEAD~N`, a branch, a tag, a changeset id, or a 64-hex commit hash; a prefix (`seq:` / `head:` / `branch:` / `tag:` / `changeset:` / `hash:`) forces which kind is meant. `swarmfile switch <branch>` is the equivalent for changing which branch the mount works on. Both are the preferred, explicit commands; `swarmfile checkout` still exists as a git-habit shim that overloads the two the way `git checkout` does — a branch name switches, a commit number / `HEAD` / `HEAD~N` restores, and the `checkout --at-seq <N>` form still works — but `switch / restore` say exactly which one you mean.

This doesn't rewrite or destroy history — it creates a new commit that matches the old state and adds it to the log. If the restore turns out to be wrong, you can restore again to undo it; nothing is ever permanently lost by restoring.

Trash and per-file rollback

Separate from the changeset system, every entry also has its own lightweight history: a per-user trash for soft-deleted files with a configurable retention window, and point-in-time rollback for individual files. It's simpler than a changeset — no multi-file atomicity, just "what did this one file look like at this timestamp." See [Trash, History & Rollback](#) for the full picture.

Where to go next

- [Branches & Merging](#) covers forking a branch to isolate risky work and merging it back.
- [CLI: swarmfile](#) has the full command reference for `commit`, `changelist`, `changesets`, `log`, `switch`, and `restore`.

Branches & Merging

Swarmfile branches follow the Perforce-streams model, not git branches. A branch is a full working line over the project's content-addressed storage, not a lightweight pointer into a shared history graph — but because storage blocks are content-addressed and shared, forking a branch doesn't copy any file data. Only the metadata diverges, and only once you actually change something. Fork a 4TB project and the fork costs you nothing until you start editing.

Branches are project-scoped

A branch applies to an entire project, not a folder inside it. You don't branch a subdirectory and leave the rest of the project on main — every mount of the project works one branch, in full, independently of every other mount. This keeps the model simple: there's never a question of which branch a given folder is "really" on.

Creating, listing, and archiving

```
swarmfile branch create look-dev-warm --from main
swarmfile branch list
swarmfile branch archive look-dev-warm
```

`branch create` forks a new branch, defaulting to a fork from **the branch this mount is currently working on** if you omit `--from` (so from inside a branch, `branch create` forks that branch — use `--from main` to fork main explicitly). `branch list` shows the active branches on the project. `branch archive` soft-archives a branch you're done with — the project's default branch (normally `main`) can't be archived.

A branch create (or a merge) is refused while a file on the source branch is still uploading, so it never silently takes that file's previous version. Retry once the upload finishes. If the uploading machine crashed, see [Troubleshooting](#).

Pruning merged branches

After a batch of branches has landed, `branch merged` lists the ones with nothing left to bring over — a branch-vs-base diff with no changed files and no conflicts (the same computation `swarmfile diff <branch>` shows, so a branch whose diff can't be computed is left out rather than assumed merged). `branch prune` archives all of them in one pass, after naming them and asking; `--yes` skips the prompt for scripts.

```
swarmfile branch merged
swarmfile branch prune
```

`prune` never archives a branch an open mount on this machine is working on — if you're checked out to a merged branch, or another mount here is on one, it's skipped (and `merged` marks it). A protected branch is listed, but archiving it is refused unless you're an org admin (exit 2). (A teammate's mount on another machine isn't visible to this check.)

"Merged" here means *merged into its own base* (the branch the diff is computed against), not specifically into `main` — the diff route only compares a branch to its base, so there's no `--into <branch>` form.

All three operations are also available from:

- **The web dashboard** — the Commits view lists, creates, and archives branches, and can trigger a merge; the **Branches** tab (see [Protecting many branches at once](#)) shows every branch's protection status plus protection rules. Both of those dashboard tabs are admin-gated.
- **The Desktop App** — a Branches panel does the same, applied immediately, with no restart required for list, create, or archive.
- **The CLI** — `swarmfile branch {list,create,switch,archive,merged,prune,protect,unprotect}`, `swarmfile protection-rule {list,create,delete,add-reviewer,remove-reviewer,list-reviewers}`, `swarmfile tag {list,create,delete}`, `swarmfile merge`, `swarmfile mr {create,list,status,approve,request-changes,review-comment,merge,ready,close,reopen,comment,comments,assign,unassign,request-review,unrequest-review,label,unlabel}`, and `swarmfile label {list,create,update,delete}`.

A typical use: fork a branch to try a look-dev variant or a design option without duplicating the whole project, then merge it back if it works out, or archive it if it doesn't.

Branching from a past commit, or a tag

By default, `branch create` forks the source branch as it is *right now*. Give it a commit ref instead (any of the forms below), and it forks the source exactly as it stood at that point — every file carrying the content it had then, in one call, instead of forking at head and then rolling the fork back with `checkout` :

```
swarmfile branch create recover --from-commit 482
swarmfile branch create recover --from-tag v1.0
```

`--from <REF>` takes the full ref vocabulary — a branch, tag, `#N / HEAD~N`, changeset id, or 64-hex commit hash — so `--from main` forks another branch, while `--from-commit` / `--from-tag` are the explicit seq/tag spellings. `--from` is optional with either: the targeted commit or tag already names which branch it lived on. If you give `--from` anyway and it disagrees with that branch, the command refuses rather than silently picking one. The fork keeps the state it started from: files inherited from an older point stay readable even after that commit ages past history retention.

This is the primitive a headless or unattended machine — a render farm node, a CI runner — uses to pin itself to an exact historical state with a single command and no Desktop App or browser interaction:

```
swarmfile commit -m "submit render job 482"
swarmfile tag create v1.0
swarmfile branch create render-fix --from-tag v1.0
```

A **tag** is a project-scoped, immutable name for a commit. It also **pins** that commit against history retention: while the tag exists, checkout/restore to it keeps working even after the commit would otherwise have been pruned, and the content it names stays stored; deleting the tag releases the pin. There's no "move a tag" command in the `swarmfile` CLI; to repoint a name, delete it and create it again (with git access, a lightweight tag can also be moved with `git push -f origin <tag>` — see [Cloning a Project with Git](#) — unless a release was published from it):

```
swarmfile tag list
swarmfile tag create v1.0
swarmfile tag delete v1.0
```

`tag create` defaults to the branch's CURRENT head — exactly what the commit above just produced, no need to look up its commit number first. Give it `--at-seq <n>` explicitly to tag an older commit instead of the head:

```
swarmfile tag create v0.9 --at-seq 482
```

The web dashboard's Commits view offers the same two actions per commit row — "Branch from here" and "Tag" — plus a tag picker next to the branch selector; the Desktop App's Branches panel has a source-type picker (branch / commit number / tag) on its create form, and its own Tags panel for list/create/delete.

Switching branches

```
swarmfile branch switch look-dev-warm
swarmfile branch switch      # omit the name to switch back to the default branch
```

Switching branches is hot — no engine restart, no interruption to the mount. The first time a machine visits a branch, switching to it pays a one-time full sync (proportional to that branch's size); a repeat visit to a branch it's already synced is a fast incremental catch-up. Switching is refused while a changelist is open — submit or cancel it first, so staged work is never silently discarded or reattached to the wrong branch. Files open in applications are closed as part of the switch, and their pending changes are saved to the branch you're leaving; an application that keeps writing to such a file afterwards gets an error, and nothing it writes lands on the new branch.

List, create, and archive don't touch the running mount and take effect immediately, same as switch now does.

The web dashboard has a separate, metadata-only branch switcher on the Files tab — a GitHub-repo-browser-style `⌵ main` dropdown that scopes which branch's file tree the dashboard is showing you. It never touches a mount: switching it there has no relationship to, and no effect on, which branch any desktop client is actually checked out to. Use it to browse another branch's files from the browser without needing a machine checked out to it.

Merging

```
swarmfile merge
swarmfile merge --record-conflicts
```

`merge` performs a whole-file three-way merge of your mount's current branch back into the branch it was forked from — its parent. That's usually main, but a branch forked from a non-main branch merges back into *that* branch. The direction is implicit either way: there's no `target` flag to set, because a branch already knows its own parent. To land a branch you aren't standing on, name it: `swarmfile merge --from feature-x merges feature-x` into its own parent without switching mounts. Because these are binary assets, there's no line-level merge to fall back on — merging is winner-take-all per file, with conflict detection when the same file changed on both sides since the fork point. By default, conflicts are surfaced as a list; `--record-conflicts` instead records conflict rows you can resolve interactively, one file at a time, rather than just being told a file collided. A recorded merge conflict belongs to the branch being merged *into*, so when you resolve it, `local` is the incoming branch's version and `remote` is the version already on the target branch: `--winner remote` keeps the target's copy. Either way the next `merge` lands cleanly, because the source branch is updated to match the kept version.

For text conflicts, the engine can try a three-way line merge before you pick a side: `swarmfile merge --diff3` batch-resolves every text-file content conflict it can, then retries the merge once. A project can make that the default by committing `.swarmfile/merge.yml` with `diff3: true` (see [Project-local files](#)); `--no-diff3` forces it off for one invocation. Anything it can't merge stays a conflict for the resolution flow below, and the web dashboard's merge button offers the same auto-resolve step when a merge comes back conflicted.

Two more flags help on big or cautious merges. `swarmfile merge --dry-run` prints exactly what would change — added, modified, deleted, renamed, moved, and conflicting files, the same read-only diff `swarmfile diff` shows — and stops without staging anything or taking the branch's merge lock. `swarmfile merge --wait` blocks until a large background apply finishes, printing progress phase by phase (the command otherwise returns as soon as the merge is accepted); it gives up after ten minutes with exit 2. And if the target branch moves after the merge was prepared, the merge refuses with `stale_merge` instead of overwriting the newer state — run `merge` again and it lands cleanly.

How this holds up on a very large project

Branching never copies file content: blocks aren't copied or re-read, only a little metadata per folder, and that metadata is written in parallel. So forking stays fast as a project grows — thousands of folders are proportionally more metadata, not more file bytes — and the first time a new branch is opened or synced, its folders are likewise read in parallel. On a 100,000-entry tree, a fork that creates 120+ directories completes in tens of seconds.

Merging is metadata-only in the same way. A merge larger than the single-pass staging budget continues as a resumable background job: the command accepts it, `swarmfile op status` shows progress while the staged diff applies, and the target head advances once, when it finishes.

Long-running work no longer depends on the terminal staying open. `merge`, `commit`, `checkout`, `revert`, `branch create / archive`, `mr merge`, `trash restore / purge`, `project delete`, `materialize`, `ignore-clean`, `hydrate free-space` and closing a mount run as **operations** in the engine: the command waits and prints the result as usual, however long it takes. Press Ctrl-C and it detaches instead of stopping — the command exits `3` and prints the operation's id — and `--no-wait` returns the id at once; `swarmfile op list / op status / op wait / op cancel` collect the outcome later, even after the engine restarts. (Ctrl-C cannot stop a commit, merge or delete part-way: those complete or fail on their own, and `op cancel` refuses them.)

Commands not on that list use their own route's socket deadline — five seconds for most, up to three minutes for a few — and if one reports `control socket request failed: no answer within ...s`, it stopped waiting but the work usually continues or is already done. Check `swarmfile branch list` (or `swarmfile mounts list`) before retrying: re-running `branch create` for an existing name answers that the branch already exists, which means the first one succeeded.

Merge Requests

`swarmfile merge` lands a branch immediately — useful when you're merging your own work and you're confident it's ready. When you want someone else to look at a branch before it lands, open a **Merge Request** instead — from the web dashboard's **Merge Requests** tab, **Merge requests** in the Desktop App's left rail, or `swarmfile mr create --title "..."`: pick the branch, give it a title, and it's posted for review.

A Merge Request shows reviewers exactly what would change — the same list of added, modified, deleted, renamed, and moved files a direct merge would land, computed fresh every time you open it, so it's never a stale snapshot. Click a changed text file to expand an inline diff right there, with syntax highlighting and word-level highlighting on modified lines; a file that conflicts with the target branch is flagged right in that list too, and clicking it shows a 3-way diff instead — your changes and the target's changes, each compared against the version both branches started from, so you can see exactly what's colliding before resolving it. Reviewers leave an **Approve** or **Request changes** verdict (optionally with a note); only the *latest* verdict from each reviewer counts, so changing your mind is just casting a new one. A verdict is also tied to the source-branch commit it was cast against: pushing a new commit afterwards retires every earlier verdict — approvals and "request changes" alike — so the approval floor has to be met again on the new head. **Merge** only becomes available once the approval gate is met and nobody's latest verdict is "request changes". On an unprotected target that gate is zero approvals; **protecting a branch** is what sets a floor, starting at one approval and configurable up to ten (a branch matched by a protection rule uses the rule's count, and when both apply the effective count is the higher of the two). There's still no required-*specific-reviewer enforcement* — raising the count (below) is the whole gating mechanism, not naming who has to be one of the approvers; requesting a specific reviewer (below) is a nudge, not a requirement — anyone's approval still counts toward the floor, requested or not. Every merge request also carries its own discussion thread — the same kind of comments you'd leave on any file — for review remarks that aren't about one specific line of the diff.

There's also a third, non-blocking **Comment** verdict (`swarmfile mr review-comment <number>`) for feedback that isn't a vote either way — it shows up in the review history same as Approve/Request changes, but never counts toward the approval floor and never blocks a merge. Unlike the other two, a Comment is allowed on your own merge request, since it can't be used to rubber-stamp your own work the way a self-approval could.

Merging a Merge Request does exactly what `swarmfile merge` does — same conflict detection, same all-or-nothing landing, same optional "record conflicts for resolution" fallback (and the same `--diff3 auto-resolve override`) if the branch drifted since anyone last looked at the diff. It also pins the source branch: if the head moved after the approval gate and CI check ran, the merge refuses with `stale_merge` rather than landing a commit the checks never saw. A Merge Request that's approved but hits a conflict on merge still 409s, same as a direct merge would; nothing lands until the branch is fixed. Closing one abandons it without merging — nothing about the branch changes — and a closed (never-merged) request can be reopened later if it turns out the work wasn't actually abandoned, from the web dashboard, the Desktop App, or `swarmfile mr reopen <number>`.

Merging also offers an opt-in **delete the source branch** checkbox (`swarmfile mr merge <number> --delete-branch`) — unchecked by default. It only takes effect once the merge itself has actually landed, and archiving is a soft **branch archive**, never a hard delete; the same protected-branch admin gate applies, so if the source branch got protected in the meantime and you're not an admin, the merge still succeeds but the branch stays as-is — you'll see a note saying so rather than the merge itself failing.

Any project member can open, review, and merge a Merge Request — it isn't admin-gated, unlike RFIs. The one restriction is that a merge request's creator can't review their own — the hub refuses with `self_review` regardless of which surface asks (web, Desktop App, or CLI); the web dashboard and Desktop App both gray out Approve/Request changes on your own merge request rather than letting you find out only after submitting.

Draft merge requests

A Merge Request doesn't have to be ready for review the moment it exists. Pass `--draft` to hold it back:

```
swarmfile mr create --title "Warmer facade material" --draft
swarmfile mr ready 3
```

A draft behaves like any other merge request for everything except merging and notifications at creation time: comments, reviews, assigning, requesting a review, and labels all work exactly the same, and assigning someone or requesting their review BY HAND still notifies them even while the request is a draft — it's only the initial "here's a new merge request" fan-out, and an **auto-reviewer's** request-review notification, that are held back at creation, on the assumption a draft isn't actually ready for anyone's general attention yet. That fan-out isn't deferred to later, either — it's skipped for good: marking a draft ready doesn't retroactively send the notification it held back at creation, so use **assignees** or **requested reviewers** if you need to be sure someone specific gets pinged. What's actually gated is merging: `mr merge` refuses a draft outright, even one that's already picked up an approval. `mr ready 3` (also available from the web dashboard and Desktop App detail view) flips it to a normal, reviewable request. There's no way back to draft once it's marked ready.

Assignees, requested reviewers, and labels

Beyond the approve/reject verdicts above, a Merge Request carries three pieces of bookkeeping — all purely organizational, none of them gating whether `mr merge` is refused:

- **Assignees** — who's on the hook to land it. One or more people, or a group.
- **Requested reviewers** — who's been asked to take a look. Distinct from an *assignee* (who lands it) and distinct from actually having *reviewed* (the Approve/Request changes verdicts above) — requesting someone is a heads-up, not a requirement, so it never changes who has to approve for `mr merge` to unlock.
- **Labels** — a per-project set of tags an admin defines once (`swarmfile label create bug --color "#FF5733"`), then anyone can attach to any merge request in that project to filter the list by.

Requested reviewers don't have to be added by hand every time. A branch protection rule can name a set of **auto-reviewers** — CODEOWNERS-style, per target branch — and every merge request opened against a branch a rule matches gets those reviewers requested automatically the moment it's created. An auto-requested reviewer is attached exactly like a manually-requested one, and only ever adds to the list — it never stops anyone else from requesting a different reviewer by hand. The one difference is notification timing on a draft: a manual request-review always notifies, but an auto-reviewer added at creation time on a **draft** MR is silently attached with no notification (same reasoning as the "new merge request" fan-out above) — they show up as a requested reviewer either way, just without the ping until the MR is marked ready and something else notifies them. See [Protecting many branches at once](#) for how to configure the list.

All three are editable from the web dashboard's detail-view sidebar, or the CLI:

```
swarmfile mr assign 3 <principal-id>
swarmfile mr unassign 3 <principal-id>
swarmfile mr request-review 3 <principal-id>
swarmfile mr unrequest-review 3 <principal-id>
swarmfile mr label 3 <label-id>
swarmfile mr unlabel 3 <label-id>
swarmfile label list
swarmfile label create bug --color "#FF5733"
swarmfile label update <label-id> --name defect --color "#000000"
swarmfile label delete <label-id>
```

Assigning, requesting a review, and attaching a label are all idempotent (doing it twice is a no-op, not an error) and don't require anything from the target — there's no accept/decline step, since these are informational, not access grants. Defining, renaming, or deleting a label is admin-gated; attaching an existing label to a merge request isn't.

From the Desktop App

Merge requests in the left rail opens a list (filterable by status, refreshing automatically as teammates act — and re-checking periodically while it's open) that drills into a detail view for one request: branch names, the changed-file list, the full review history, the discussion thread, and a review composer — click through to Merge or Close from there, both behind the same confirm step every other consequential Desktop App action uses. Same as the web dashboard, click a changed text file for its inline diff, or a conflicting file for a 3-way diff. A **+ New Merge Request** button on the list opens a small inline form (source branch, title, optional description, and a "Create as draft" checkbox) rather than needing the CLI or the browser just to propose a review.

A **draft** shows a **Draft** badge in place of the usual approval pill, and its Merge button is replaced by **Mark ready for review** until you flip it. If a target branch [requires CI checks to pass](#), a merge blocked purely on a failing/missing check says so specifically ("Required checks haven't passed yet") rather than the generic "needs approval" message. Assigning, requesting a review, and labeling are done from the web dashboard or the CLI — the Desktop App shows them read-only.

From the CLI

```
swarmfile mr create --title "Warmer facade material"
swarmfile mr list
swarmfile mr status 3
swarmfile mr approve 3
swarmfile mr request-changes 3 --note "needs another pass"
swarmfile mr merge 3
swarmfile mr close 3
swarmfile mr reopen 3
swarmfile mr comment 3 --message "Looks good once the fascia texture is swapped."
swarmfile mr comments 3
```

`mr create 's` `--from` defaults to whatever branch the mount is currently on, so a CI job that forked its own branch and did the work can open a review with no separate branch-name bookkeeping. `mr status <number>` doubles as a CI gate — its exit code is `0` if approved, `2` if open but not yet approved, `1` on any other error. `mr comment / mr comments` post to and list the same discussion thread the web dashboard and Desktop App show.

`mr approve / request-changes / merge / close / reopen` put review actions a terminal command away, not just the web dashboard or Desktop App. The hub's `self_review` gate stops the merge request's own creator from reviewing it; an approval from a second credential counts, so treat review-capable credentials like any other approval authority in your org. If you want an independent-approval guarantee, protect the target branch with `--required-approvals` set above 1 ([below](#)).

Protecting a branch

`swarmfile merge` and a direct `mr merge` both bypass review the moment nobody insists otherwise. **Protecting** a branch removes that choice for changes arriving as a *merge*: once set, a bare `swarmfile merge` into it (or `mr merge` on a request that hasn't cleared the branch's approval/checks gates) is refused outright with `protected_branch` — the branch's changes have to arrive through an approved Merge Request.

Protection governs **what lands on** the branch, whether it arrives as a merge or a direct write. Once set, a bare `swarmfile merge` (or `mr merge` on a request that hasn't cleared the branch's gates) is refused, and so is every direct write from a mount: saving, editing, renaming, moving, or deleting a file, creating one, staging/committing directly, rolling back, resolving a conflict, or reflecting LFS content all answer `409 protected_branch`. The way through is a branch plus an approved Merge Request; project ACLs still decide *who may ask*, but they cannot make a protected branch writable.

Before you protect: update every engine first. Protection is enforced by the hub, and the desktop app on each machine decides how to react to a `409 protected_branch` refusal. Update the app on every machine that works on a branch **before** you protect it, so staged work is held for you to move to a branch rather than discarded.

```
swarmfile branch protect release
swarmfile branch protect release --required-approvals 2
swarmfile branch unprotect release
```

Setting or clearing protection is admin-gated — a plain project member gets refused (the CLI exits `2`, distinguishable from a hard error), same as trying to archive a protected branch. Every protect/unprotect leaves an audit trail: who did it, and when, shows up in the dashboard's Activity feed (owner-visible, same as an ACL-mode change). The web dashboard's Files-tab branch switcher and the Commits toolbar both offer a protect/unprotect toggle for admins, alongside a  badge everyone sees; the Desktop App's Branches panel has the same toggle.

On a protected branch the mergeability rule is " $\geq$ the branch's effective required-approvals (1-10, default 1), 0 requested-changes"; on an unprotected branch the approval half is 0, so a direct `mr merge` needs none until the branch is protected. A single `changes_requested` verdict still blocks a merge outright no matter how many approvals exist, and that part isn't configurable. `--required-approvals` (1-10) sets the approval half for this one branch; leaving it off keeps the default of 1 while the flag is set. There's still no required-specific-reviewer enforcement — raising the count is the whole gating mechanism, not naming who has to be one of the approvers; [requesting a specific reviewer](#) flags whose attention it needs without making their approval individually required.

```
swarmfile protection-rule create release --require-checks-to-pass
```

`--require-checks-to-pass` exists on **protection rules only**, not on `branch protect` — to gate one exact branch on CI, use a no-wildcard rule like the one above. It adds a second, independent gate alongside approvals: once set, a merge is refused unless the *latest* run of every check reported at the merge request's exact current head — not an earlier commit it happened to pass on — is green. A rule that last reported at an earlier commit isn't carried forward as still-required, but if no check has reported at the head at all, the gate fails closed rather than passing vacuously — the same "no news is bad news" convention GitHub's required-status-checks uses, so the gate can't be quietly satisfied by a check nobody's runner has gotten around to yet. This stacks with `--required-approvals` rather than replacing it — both have to clear — and the merge request detail view exposes checks state separately from approval state, so an MR can visibly be "approved, blocked on checks." A check doesn't have to come from Swarmfile's own [runner](#) — any CI system that can call an authenticated API can report one.

Protecting many branches at once

Protecting one branch by name is fine for `main`, but not for a `release/*` line that grows every sprint. A **protection rule** protects every branch whose name matches a glob pattern — including ones forked *after* the rule was created, with nothing to re-apply:

```
swarmfile protection-rule create "release/*" --required-approvals 2
swarmfile protection-rule list
swarmfile protection-rule delete "release/*"
```

`*` matches any run of characters, `?` matches exactly one; a pattern with no wildcards (`protection-rule create main`) is just another way to protect one exact branch — equivalent to `branch protect main`, not a second, different mechanism. A branch counts as protected if its own flag is set OR any rule matches its name, so deleting a rule can un-protect a branch that was never explicitly protected itself. `--required-approvals` composes the same way: the effective count for a branch is the MAX of its own value (if flag-protected) and every matching rule's — the same "most restrictive wins" resolution the boolean uses, just with the higher number instead of `true`. Same admin gate, same exit-2 convention, and same audit trail as the per-branch toggle — a rule is arguably the more consequential of the two, since one rule change can flip many branches (including future ones) at once.

The two mechanisms don't override each other: `branch unprotect release/1.0` only ever clears that branch's OWN flag. If a rule like `release/*` still matches it, it stays protected regardless, and every surface says so — the CLI/Desktop App/web response reports the branch's real, effective state after the unprotect, not just "done."

A protection rule is also where [auto-reviewers](#) live:

```
swarmfile protection-rule add-reviewer "release/*" <principal-id>
swarmfile protection-rule remove-reviewer "release/*" <principal-id>
swarmfile protection-rule list-reviewers "release/*"
```

Every merge request opened against a branch matching `release/*` requests a review from each configured principal automatically, the moment it's created — no one has to remember to add them by hand. The list is matched against the merge request's *target* branch, same as the rule's own pattern; it's additive to whatever reviewers get requested manually, never a replacement.

Rules are managed from the CLI or the web dashboard's **Branches** tab (admin-gated, alongside Commits) — a single view of every branch's protection status plus every active rule, since neither the Files-tab switcher nor the Commits toolbar shows more than one branch at a time. The same **Branches** tab has an **Edit** button on any already-protected branch to raise its required-approval count without unprotecting it first; the Desktop App's Branches panel has the same Edit button on its own protected-branch rows, an inline number field replacing the row's usual buttons rather than a separate dialog. Protection rules themselves — creating or deleting a `release/*`-style pattern — are CLI/web-only; a Desktop App user sees a rule-protected branch's  badge (with an `×N` suffix once the count is above 1) correctly every way, since the hub resolves the effective protection state and count server-side before any client sees them, just can't create or delete a rule from the Desktop App.

Protecting a branch — by flag or by rule — is only as strong as the Merge Request review it forces you through, and it inherits the caveat above: the hub stops an MR's own creator from approving it, while an approval from a second credential counts. `--required-approvals` above 1 raises the bar — clearing a 2-reviewer rule takes two approval credentials, not one — but no server-side check can prove that two accounts are two different humans without identity verification. If you protect `main` specifically to guarantee a human looked at every merge into it before it landed, keep review-capable credentials out of any automation that opens merge requests.

A different kind of conflict

Branch merges aren't the only place conflicts happen — and the other kind is more common. If you and a teammate both edit the same file on the *same* branch while one of you was offline, Swarmfile catches it when the offline edit tries to land: the change was made against a base version that's no longer current. Rather than

silently keeping both copies and leaving you to notice, it records a conflict and **refuses further writes to that file** until it's resolved. That's why the file shows as read-only on macOS and Linux and an application's save fails with "permission denied" ("access denied" on Windows), and why `swarmfile conflicts` exists — a filesystem has no way to say "someone else changed this too." The web file browser badges the row with **Conflict** so it's visible without opening the file, the Desktop App's **Conflicts** panel lists it with its path (and the branch, when it isn't the one this mount is on), and the CLI is where a terminal sees the reason.

Resolving is a deliberate choice between several outcomes. In the web dashboard and the Desktop App (both open the same picker):

- **Keep mine** — your offline version becomes the current head; the other edit is superseded.
- **Keep theirs** — keep the version already on the hub. Your offline edit isn't discarded: it's saved in the file's **history** (as a version owned by you), so you can restore it later with a [rollback](#) or from the History view. Choose **Keep both** instead when you want your version visible alongside the file right away.
- **Keep both** — the hub keeps the current head, and your version lands alongside it as a sibling file with a conflict suffix, so nothing is lost and you reconcile them by hand.
- **Try an automatic merge** — for file types with a registered merge tool (currently plain text and common source/config formats), click "Try automatic merge" and Swarmfile runs a three-way line merge of your edit against theirs. If it succeeds cleanly, a fourth choice appears — the merged result — for you to review and accept instead of picking a side. If the file type isn't supported, or both edits touched the same lines, the attempt just fails harmlessly and you're left with the three choices above.

From the CLI, `swarmfile conflicts resolve <path>` covers the first three non-interactively via `--winner local|remote|keep-both`. `--auto` runs the same built-in three-way merge the picker's "Try automatic merge" does — no external tool and no git config needed, and it exits 2 when a real conflict remains (so a script can tell "a human must choose" apart from "the engine was unreachable" and fall back); add `--all` to sweep every conflicted file on this mount's branch in one pass (rows on other branches are counted and skipped unless `--include-other-branches` is given — pass it or `merge --diff3` when landing a merge, whose conflict is recorded under the target branch). `--tool [name]` is the option the picker doesn't have: it hands the conflict to an **external mergetool** — `kdiff3`, `Beyond Compare`, VS Code's merge editor, anything speaking git's own `$BASE / $LOCAL / $REMOTE / $MERGED` protocol — auto-discovered from your git config. Unlike the built-in merge, `--tool` works on any file type, binary included, since the tool does the actual reconciling rather than Swarmfile attempting a line-based diff. See [CLI: swarmfile](#) for the full flag reference. `--tool` is CLI-only by design: an external tool needs a real terminal/GUI session to run in, which only a foreground CLI process has.

Like branch merging, whole-file is always available as a fallback here — a binary asset never has a line-level merge to fall back on, which is why `keep-mine/keep-theirs/keep-both` always work regardless of file type. Until you resolve the conflict, the file stays read-only; once you do, writes resume with no restart.

An unresolved conflict also stops every other operation that would move that file's content: a branch merge, a point-in-time restore, a commit revert, and a rollback all refuse with `conflict_unresolved` until it's resolved. Resolve it first, then re-run the operation.

Where to go next

- [Version Control](#) covers changesets, checkout, and per-file history — the model branches sit on top of.
- [CLI: swarmfile](#) has the full command reference for `branch` and `merge`.

Trash, History & Rollback

Every change to a file is timestamped, and you can always roll back to an earlier point in time. This page covers the lightweight, per-entry mechanism for that: trash and per-file history. For the multi-file, atomic changeset-and-checkout system, see [Version Control](#) — that's a different, heavier tool for coordinating a whole set of changes together, not a replacement for this one.

Trash

Deleting a file doesn't remove it immediately. It soft-deletes into a per-user trash, visible under the project's **Trash** tab in the web dashboard (open the project, then Trash in its tab strip). There is no cross-project trash view — each project's trash is its own.

You can delete from either place — your file manager (Finder, Explorer, or the command line, on the mounted drive) or the web dashboard. Both do the same thing: the file goes to trash, not away. Deleting a folder takes its contents with it. The dashboard does that in one step; your file manager removes the files inside one at a time and the folder last.

Restoring a folder brings back the folder and everything that was deleted with it in the same step. A folder deleted from your file manager went one file at a time, though, so when you restore it, the web dashboard lists the files that were deleted inside it just before it and asks whether to restore them too. It asks rather than restoring them automatically, because a file you deleted on purpose a moment earlier looks the same.

Retention is per-user, not just per-org: your own account setting, if you've set one, wins; otherwise it falls back to the org's default; otherwise a global 90-day default applies. Set your own under **Settings → Account → Trash & history retention** (1 to 3650 days), or clear it there to go back to the organization's window. An org's default retention window is settable anywhere from 1 to 3650 days and controls how long a deleted file stays recoverable before it's permanently purged, for anyone in the org who hasn't set their own override. Until that window elapses, restoring a file from trash brings it back in place, at its original path. Retention windows are applied by a daily cleanup, so a change takes effect at the next daily pass rather than instantly.

From the CLI, `swarmfile trash list` shows your own soft-deleted entries, `swarmfile trash restore <entry-id>` brings one back (add `--with-nearby` to also bring back what was deleted inside a folder just before it), and `swarmfile trash purge <entry-id>` deletes it permanently — "Delete Forever," ahead of the retention window, not recoverable. `--preview` on `purge` reports how many descendants a folder would take with it before you commit to removing anything. Because it can't be undone, `purge` asks for confirmation first; `--yes` skips the prompt (a non-interactive shell or `--json` needs it). A very large folder restore or purge runs as a durable background job rather than an inline cascade: the command waits for it by default, and Ctrl-C or `--no-wait` detaches with exit 3, followed with `swarmfile op status / op wait` ([long operations](#)). `op status` shows the job's `done / total` progress while it runs, and in the dashboard the same operation shows progress on the row's own button — a restored folder reappears at once, with its contents filling in behind it.

History and rollback

Because every change is timestamped, a file's full edit history is available at any time, and you can roll back to any point in that history — not just the most recent save. This is scoped to a single entry: rolling back a file doesn't touch anything else in the project.

Every saved version vs. named commits

New projects start on **Keep every saved version**: each automatic save stays restorable, so any point in a file's history can be brought back — not just its named commits. If you'd rather keep storage leaner — on a busy, high-save-volume project, say — choose **Named commits only** when creating the project, or switch later in the project's settings (or with `swarmfile project set-keep-full-history false`): individual save versions then expire after the retention window, and committed changesets are kept for that same window — tag one to keep it indefinitely (Keep every saved version keeps everything). The setting is per project, not per file; an org owner can change which choice new projects start with under **Settings → Organization** (which applies only to projects created afterwards).

When history is kept for less time

On a project set to **Named commits only**, a file's history is kept for your retention window, with one exception: if a project's stored file information grows very large (millions of files and versions), Swarmfile automatically keeps history for less time — never below a minimum of 7 days by default — so the project stays under its storage limit. It returns to the full window once the project is comfortably below that point again. Org owners and admins can see the window currently in force, and why, on the project's **Project settings** tab. See [Project storage limit](#).

Trash and history vs. version control

Trash and history operate one entry at a time, with no concept of grouping changes across files. If you need to restore a whole folder or project to how it looked before a coordinated set of changes — a batch of asset updates that need to move together, for example — that's what the changeset/checkout system in [Version Control](#) is for.

Verifiable History

Version history is only as trustworthy as whoever keeps it. Swarmfile's history is built so you don't have to take our word for it: every commit is **hash-chained** to the ones before it, the service **signs checkpoints** of each branch with a published key, and a standalone tool, `swarmfile-verify-history`, lets you check a project's history yourself — on an audit host or in CI, with no Desktop App or mount.

This page explains what that gives you and where its limits are. For flags and exit codes, see the [CLI reference](#).

Hash-chained commits

Every commit has a **commit hash** computed from its own facts: the content-addressed id of the project tree it records (which in turn is built from the ids of every file's contents) and the hashes of its parent commits. Because each commit's hash covers its parents' hashes, a commit's hash covers the whole history behind it. Change one file in one old commit, or remove a commit from the middle, and every hash after it stops matching.

Commits get their hash once they're **confirmed**. Swarmfile doesn't compute commit hashes on the server: a client derives the commit (the Desktop App does this after a commit, or `swarmfile git index` on demand), and a second, independent derivation confirms it — a teammate's app, a fetch, or a project API key used in CI; an owner can also allow their own app to self-confirm. Commits made where no app runs (a web upload, a revert, a merge) are derived by Swarmfile's server, which counts as one independent derivation: the first matching derivation from a client confirms them. Commits pushed with git work the other way round: the push derives them, and the server's own derivation a couple of minutes later confirms them. Self-confirm covers only the owner's own app, so on a self-confirm project these commits stay provisional until a client derives them (a fetch is enough). Until then a commit is *provisional*: it's in your history and fully usable, it just isn't part of the hash chain yet. `swarmfile git status` reports how many commits on a branch are still unconfirmed.

Signed checkpoints

Hash chaining proves the history is internally consistent. It doesn't, on its own, stop someone who controls the server from replacing the whole chain with a different, equally consistent one. That's what checkpoints are for.

Periodically (hourly, for every branch that has changed since its last checkpoint) the Swarmfile service signs the branch's latest **confirmed** commit hash with an **Ed25519** key. The verifier requires every checkpoint published for a branch to land on the chain it walks, so a history that has lost or replaced a checkpointed commit fails verification. The signing key can be rotated; the service publishes its current and retired public keys so older checkpoints stay verifiable. Verifiable history is offered where the service has checkpoint signing enabled, and the verifier reports it as unavailable where it isn't (see [What it proves, and what it doesn't](#)).

Checking it yourself

`swarmfile-verify-history` is installed with every Desktop App and runs on its own:

```
swarmfile-verify-history \
  --hub-url https://hub.swarmfile.com/orgs/<org-id>/projects/<project-id> \
  --org-id <org-id> \
  --project-id <project-id> \
  --branch main \
  --header "Authorization: Bearer $SWARMFILE_API_KEY" \
  --pin-key ./swarmfile-history-key.pub
```

It fetches the branch head, then walks back commit by commit to the start of history (or to a commit you name with `--checkpoint-hash`), recomputing every hash and checking every parent link. It then fetches the branch's checkpoints and checks each signature. Any mismatch is a failure with a non-zero exit, so you can gate a release pipeline on it; `--json` gives a machine-readable report.

Pin the key. Without `--pin-key`, the verifier checks signatures against the public key the same service is serving you, which proves the checkpoints match *that* key but not that the key is genuine. Pass the key you recorded out-of-band (from an earlier run, or from us directly) with `--pin-key` and the run only accepts signatures from keys you supplied. That's the mode that gives you independent assurance. Repeat `--pin-key` to accept a rotated key alongside the old one.

What it proves, and what it doesn't

- **It proves** that the history the service serves you is internally consistent — no edited commit, no broken link — and, with a pinned key, that the checkpoints on it were signed by the key you trust.
- **It doesn't, by itself, prove** that this is the same history you saw last month. The checks run against the checkpoints the service lists, so keep your own record of a commit hash you care about (a release, a delivery) — or of the checkpoints from an earlier run — and confirm it's still on the branch.
- **Provisional commits aren't covered** until they're confirmed. A branch with no confirmed commits yet reports that, rather than passing.
- **It needs the service to be reachable.** The verifier needs no Desktop App or mount, but it reads the commit objects and checkpoints from the service as it walks; it doesn't verify an exported bundle offline.
- **Checkpoints require checkpoint signing on the service.** Where the service doesn't have it enabled, the verifier reports verifiable history as unavailable — an informational result unless you pinned a key, in which case it's a failure.

Where you'll see it

- **Releases** of a [public project](#) marked **commit-pinned** name the exact commit hash they were published from. The badge is the pointer; `swarmfile-verify-history` is how you check it.
- **git clone** of a project uses the same confirmed commits: each project commit in the clone carries a `Swarmfile-Commit:` trailer naming the Swarmfile commit it came from (the graft commit that anchors history at the cut, if any, has no trailer). See [Cloning a Project with Git](#).

- **Reproducible checkouts** — `swarmfile materialize hash:<commit-hash>` writes exactly that commit's tree, the most precise pin you can give a CI job (tags can be moved; hashes can't).

Git and Swarmfile

For an organization, Swarmfile is a version-control system in its own right — and for large, binary-heavy projects a more flexible one than Git: there is nothing to clone or keep in sync (the project streams from a mounted drive), access is per folder, and locking, review, rollback and releases are built in. Nothing here asks you to leave Git: the three paths below are where the two meet, so you can adopt Swarmfile self-serve through whichever fits and move more of your work over at your own pace.

Swarmfile isn't a full code-hosting forge — there's no issue tracker, and code review happens through Swarmfile's own merge requests — but a project with git access enabled serves a real git repository you can `git clone swarmfile://` and `git push` to. Plenty of teams that use it live in git too, so there are three separate ways the two fit together. They solve different problems, and you can use any combination:

	What it's for	Where your source lives
Swarmfile as your git-LFS server	Keep your repo on your git host; send the big binary files to Swarmfile.	Your git host
<code>git clone swarmfile://</code>	Get a real git repository of a Swarmfile project, for git tooling, CI, or an offline archive, and push commits back.	Swarmfile (the mounted drive)
The <code>swarmfile</code> CLI's git-style commands	Do version control on the mounted drive with the verbs you already know.	Swarmfile (the mounted drive)

1. Swarmfile as your git-LFS server

Point git-LFS at a Swarmfile project, and `git lfs push / pull` move your large files through Swarmfile while everything else stays in your existing git host. You get storage, deduplication and quota without bloating the host, pushes up to 256 GiB per object with the transfer agent, `git lfs lock` that is the same lock as a lock on the drive, and an exit with no lock-in: `git lfs fetch --all` with stock git-LFS pulls everything back. Pushed objects can also be reflected onto the mounted drive as ordinary files (through the transfer agent — a stock `git-lfs push` stores the object but doesn't place it on the drive).

git-LFS is the one surface where a managed project can hold plaintext by design: objects a stock `git-lfs` client pushes, and the desktop agent's very large uploads, are stored unencrypted. Anything that must be encrypted at rest belongs on the mounted drive; git-LFS isn't available on end-to-end-encrypted projects.

Start with [Git-LFS](#).

2. `git clone` of a Swarmfile project

Once an owner turns on git access for a project, anyone who can read all of it can run:

```
git clone swarmfile://<org>/<project>
```

and get every branch and tag as real git history, with large files as LFS pointers served by Swarmfile. `git fetch` brings it up to date, and `git push` sends new commits back as Swarmfile commits with their git ids unchanged. Pushes are fast-forward only, a push to a protected branch opens a merge request instead, an existing repository can be pushed whole into a new project, and a pushed commit must look exactly like what a clone would produce (large files as LFS pointers, no submodules; executable bits are kept) — the [guide](#) has the rules.

A pushed commit is *provisional* until a second, independent derivation confirms it. Swarmfile's server does that on its own a couple of minutes after the push; a teammate's fetch also confirms it, and `swarmfile git index` confirms a whole branch's unconfirmed commits in one run. A clone can opt out with `git config swarmfile.confirm false`. See [How a pushed commit gets confirmed](#).

Start with [Cloning a Project with Git](#).

3. The `swarmfile` CLI speaks git

On the mounted drive, the `swarmfile` CLI uses git's vocabulary where the meaning carries over — and where it doesn't, it tells you instead of failing.

Commands that work the way you expect: `commit` (including `--amend`, `message-only`), `log`, `show`, `diff`, `branch`, `switch`, `checkout`, `restore`, `revert`, `merge`, `tag`, `describe` and more. Your `~/gitconfig` aliases work too (`swarmfile co`, `swarmfile st`). [Coming from Git](#) has the full cheat sheet.

Git habits that run the real thing:

You type	It runs
<code>swarmfile add</code>	Opens a changelist (if one isn't open), so your next saves are staged into it.
<code>swarmfile stash / stash push / stash save</code>	Parks the open changelist (<code>changelist park</code>).
<code>swarmfile stash list / stash show</code>	Lists parked work (<code>changelist parked</code>).
<code>swarmfile stash pop / stash apply</code>	Resumes parked work (<code>changelist resume</code>).
<code>swarmfile stash drop / stash clear</code>	Discards parked work (<code>changelist discard</code>), asking first.
<code>swarmfile reset --hard <rev></code>	<code>swarmfile restore <rev></code> — one undoable commit, not a history rewrite.
<code>swarmfile clone <project-id> <path></code>	Opens a mount of that project at that path (<code>-b <branch></code> pins a branch; <code>--checkout</code> writes a plain directory instead).

Git habits that print guidance instead: `push`, `pull`, `clone` (with anything other than a project id and path), `rebase`, `cherry-pick`, `rm`, `mv`, `cp`, `mkdir`, `touch`, `worktree`, `grep`, `init`, `remote`, `clean`, `reflog`, `mergetool`, `difftool`, `shortlog`, `submodule`, `bisect`, `gc`, `prune`, `fsck`, `archive` and `notes` each print a one-line reason and what to do instead — `push` and `pull` explain that saves already sync, `worktree` points at `swarmfile` mounts open, `grep` at `swarmfile-search`, `mergetool` at `swarmfile conflicts resolve --tool`. These work even when the Desktop App isn't running. The [CLI reference](#) has the details.

Compatibility at a glance

You want to...	Today
Keep source in your git host and store big files in Swarmfile	Works — git-LFS server
<code>git lfs lock</code> files, including files only in git	Works
Leave with all your LFS objects using stock git-LFS	Works — <code>git lfs fetch --all</code>
<code>git clone</code> a Swarmfile project	Works , once git access is on for the project — guide
<code>git fetch</code> / <code>git pull</code> new commits into that clone	Works , incrementally
<code>git clone --depth N</code> (shallow)	Works , once the project's commits are indexed; <code>--shallow-since</code> / <code>--shallow-exclude</code> don't
<code>git push</code> to a Swarmfile project	Works , fast-forward only, on managed and unencrypted projects — rules : deleting a branch archives it, lightweight tags become Swarmfile tags (<code>--force</code> moves one and <code>:refs/tags/<tag></code> deletes one, unless a release was published from it), a push to a protected branch opens a merge request, an existing repository can be pushed into a new project in batches of up to 100 commits, a merge brings unpushed side history along on an auto-archived <code>merged/<sha></code> branch, and a commit changing more than 5,000 paths goes through a chunked, resumable push session with rollback; force pushes to a branch and annotated tags are refused — <code>swarmfile git force-forward</code> lands a force's end state as a forward commit instead
Import an existing git repository into a new Swarmfile project	Works — <code>swarmfile git import</code> creates the project (git access on, full history) and pushes every representable branch and tag with their upstream ids; annotated tags become lightweight (<code>--annotated-tags skip / fail</code> to opt out); a default branch that uses Git-LFS, submodules or an octopus merge is refused, while the same problem on another branch skips and lists it; re-running the command syncs upstream changes into the project it made (<code>--new</code> for a second project)
Partial clone (<code>--filter</code>)	Not supported — large files already arrive as LFS pointers
Clone over an HTTPS URL	Not supported — use <code>swarmfile://</code>
Clone an end-to-end-encrypted project	Not supported — mount it, or use <code>swarmfile materialize</code>
Fetch an arbitrary commit by id	Not supported — branches and tags only
Executable bit in a clone or push	Works — executable files are mode <code>100755</code> both ways (<code>swarmfile chmod</code> sets it on Windows)
Empty folders in a clone	Not carried — git can't store them
Submodules	Not supported in Swarmfile's own version control — <code>swarmfile submodule</code> prints guidance
Rebase, cherry-pick, history rewriting	Not supported by design — Swarmfile never rewrites commits; <code>commit --amend</code> changes only the message

Which one should I use?

- **Your source code already lives in git and you're happy with that** → the [git-LFS server](#) on-ramp. Nothing about your workflow changes except where big files go — and when you want the drive, review and hosted CI for the whole project, turn on git access and the same project serves the code too.
- **Your repo lives in a git host and you want it as a Swarmfile project** → `swarmfile git import <url>` brings the whole history, branches and tags over in one step ([reference](#)).
- **Your project lives on the Swarmfile drive, and some tool needs a git repository** → `git clone swarmfile://` .
- **You're doing version control on the drive and your fingers type git** → the CLI, with [Coming from Git](#) open in a tab.
- **CI that builds from a branch** → you may not need a clone at all: [Hosted CI](#) runs `.swarmfile/ci.yml` workflows on Swarmfile's containers against the branch directly, the self-hosted [Runner \(Headless CI\)](#) runs jobs on your own machine, and `swarmfile materialize` writes any ref to a plain directory.

Cloning a Project with Git

A Swarmfile project is normally something you **mount**, not clone: the drive is live, saves sync as you make them, and there's no working copy to keep up to date. But some tools want a real git repository — a build system, a code-review or static-analysis tool, a CI job, an archive you keep offline. For those, an owner can turn on **git access** for a project, and anyone who can read the whole project can then:

```
git clone swarmfile://<org>/<project>
```

The result is an ordinary git repository: real commits, branches and tags, with large files as git-LFS pointers that Swarmfile serves. `git fetch` brings the clone up to date, and `git push` sends your commits back as Swarmfile commits (fast-forward only — see [Pushing](#)), and an existing repository can be pushed into a new project. Starting from a repository that already exists elsewhere? `swarmfile git import` creates the project from it in one command.

New to how git and Swarmfile fit together? [Git and Swarmfile](#) compares this with the git-LFS server and the CLI's git-style commands.

Turning on git access

Git access is a per-project setting that covers every branch and tag. A project owner or org admin turns it on:

1. In the web dashboard, open **Projects**, use the project's menu, and choose **Git access....**
2. Read the warning, tick **I understand this cannot be turned off or re-tuned**, and select **Enable git access**.

The desktop app has the same switch: with the project mounted, open **Project settings** → **Git access** → **Enable git access....** Or from the terminal: run `swarmfile git enable`. Both freeze the same record and print what was frozen.

Git access is off until an owner enables it. A member who tries to clone, push, or index a project without it gets the refusal, and the hub nudges the org's owners once (`git_enable_requested`) so the request doesn't go silent.

It's a one-way switch. At the moment you enable it, Swarmfile fixes the rules it uses to turn the project into git: the conversion version, the size above which a file becomes an LFS pointer (1 MiB by default), and where history starts. Those rules never change afterward, because changing any of them would change every commit id a clone has already seen. If a different large-file boundary matters for the repo, set it in the same breath as enabling — `swarmfile git enable --threshold-bytes <N>` freezes that value instead of the 1 MiB default. Once enabled, the dialog shows the frozen settings instead of the switch, plus the exact `git clone` command to copy. That command names the project by its id — the form that also works on a CI machine signing in with a project API key.

Before you enable it, two things must be true, and the hub refuses the enable (and says which) if they aren't:

- **The project keeps every saved version.** Git access needs the project's full history to stay restorable. Turn it on with **Keep every saved version** in the project's history settings, or `swarmfile project set-keep-full-history true` (see [Trash, History & Rollback](#)).
- **Every large file has a recorded id.** A file at or above the LFS threshold reaches git as an LFS pointer, and a pointer needs the file's SHA-256. Files saved through the Desktop App get one as they're saved; older files may not. `swarmfile git status` shows how many are missing, and `swarmfile git enable` records any that are missing automatically before it freezes the view (run `swarmfile git backfill-oids` ahead of time if you want to see the work finish first).

The enable is also refused on an **end-to-end-encrypted** project (git access isn't available for E2E projects; mount the project, or use `swarmfile materialize`, instead), when a branch or tag head can't be read, and when the project's history is too long to check in one pass.

Where history starts. When you enable git access, Swarmfile walks back from every branch and tag. If some older commit's content is no longer stored (for example, history from before the project kept every saved version), history is cut there: that commit appears in git as a commit with no parent. Everything from that point forward is complete.

Cloning

```
git clone swarmfile://acme/website      # by org and project slug
git clone swarmfile://acme/website site # into a directory of your choice
swarmfile git clone acme/website site    # the same, through the CLI
```

The forms the URL takes:

URL	Meaning
<code>swarmfile://<org>/<project></code>	The project on the Swarmfile service your Desktop App is signed in to (or <code>SWARMFILE_HUB_URL</code>). Slugs or ids both work.
<code>swarmfile://<hub-host>/<org>/<project></code>	The same project on a specific Swarmfile service, always over HTTPS.
<code>swarmfile://<project-id></code>	A project by its id, in your configured organization.

`swarmfile git url <project>` prints the URL for a project, handy for `git remote add origin "$(swarmfile git url acme/website)".` `<project>` can be `<org>/<project>`, a project id, or a bare project name — a bare name is qualified with your configured organization, and the printed URL then works anywhere that can reach the project.

In a URL, each segment is the org's or project's **web address** (the lowercase slug, e.g. `yadc` for a project named YADC) or its id. A slug is matched ignoring case; when no address matches, a display name that is unique among your orgs (for the org segment) or the org's projects (for the project segment) also resolves — a name matching more than one is refused with the candidates listed, rather than guessed.

What you need installed. Cloning uses a small git *remote helper*, `git-remote-swarmfile`, which every Desktop App installer ships (on macOS the installer also links it into `/usr/local/bin`; on Linux it's in `/usr/bin`). No running drive is needed. If `git` says it can't find a helper for `swarmfile`, use `swarmfile git clone`, which finds the helper next to the `swarmfile` binary even when it isn't on your `PATH`. Git 2.11 or later is required.

There's no branch in the URL: a clone fetches every branch and tag, with the project's default branch as HEAD — `main` unless its default branch is something else: a different [initial branch](#) at creation, or the branch an [import](#) brought with it. Use git's own options to narrow it, for example `git clone -b design-b --single-branch swarmfile://acme/website`.

On a [protected project](#), your **current** access governs what git can materialize: a teammate added to the project can clone its whole history, including commits made before they joined — the derived ids depend on tree content, and the hub checks permissions as they are now. The hub separately remembers what access existed at each commit's time, which is what keeps point-in-time history views faithful; that replay isn't what git reads. If your access is revoked, git reads stop too, just like the file browser.

Signing in

The helper looks for a credential in this order:

1. **SWARMFILE_API_KEY**. A project API key in the environment wins over everything else, including a Desktop App session, and the helper says so when both are present.
2. **Your Desktop App sign-in**. On a machine where you're signed in to the Desktop App and no explicit key is set, `git clone` just works.
3. **A project API key from git's credential store**. On a machine with no Desktop App session (a CI box, a container), store a project API key (`sf_key_...`) as the password for the Swarmfile host with `git credential approve`, and the helper picks it up when no explicit key or session is present.

A key is project-scoped, so for a **private** project use its **id** in the URL — `swarmfile://<org>/<project-id>`. Resolving a project *slug* needs a session (a key can only resolve a public project by slug); the refusal says to use the project's id. A stored credential (option 3) is never silently preferred over an explicit `SWARMFILE_API_KEY`.

For a scripted sign-in, `SWARMFILE_OIDC_REFRESH_TOKEN` works and is **not** one-shot: the IdP rotates refresh tokens on use, but the live successor is persisted under the Swarmfile cache directory, and later commands — the push after a clone, a restarted engine, any other git command — reuse that stored successor instead of the spent environment value. Supplying a *different* token later still wins. A stored Desktop App session or a project API key remains the simplest CI credential (no cache-dir file to carry between steps).

A clone needs to read the **whole** of every commit it fetches, and the hub checks permissions against the access you hold **today**. If a folder is restricted for you now, it can't hand those commits over complete, and a clone or fetch that has to walk one is refused with a message saying so rather than producing a repository with silent holes — grant read and the same clone succeeds, including commits made while you were restricted. (The hub also records what access existed when each commit was made; that replay answers explicit point-in-time history views, not git.) The mounted drive and [swarmfile materialize](#) follow your access today in the same way. Use an owner's or a fully-permitted member's credential, or a project API key, for clones of restricted content.

What a clone contains

- **Branches and tags**. Every non-empty branch becomes a git branch; tags become lightweight git tags. An empty project clones as an empty repository.
- **Commits**. Each Swarmfile commit becomes a git commit, with your message and a `Swarmfile-Commit:` trailer naming the Swarmfile commit it came from. A commit with no message reads `(no message)`. Authors are recorded by account id; the helper writes a local mailmap so `git log` shows the names of people in your organization.
- **Unsaved-to-a-commit work**. If a branch has saves since its last commit (autosaves that haven't been rolled into one yet), the clone ends that branch with one extra commit, authored by `swarmfile` with the message "swarmfile snapshot of the live tree", so what you clone matches what's on the drive.
- **Large files as git-LFS pointers**. Files of 1 MiB or more come through as LFS pointers, with a generated `.gitattributes` in each folder that has them (added after any `.gitattributes` you already had). The clone is configured to fetch LFS content from Swarmfile automatically — see [Git-LFS](#) for the transfer agent and performance notes.
- **Symlinks** stay symlinks. **Empty folders** aren't carried (git can't store them). **Executable files** (marked with `chmod +x` on the macOS or Linux drive, or [swarmfile chmod +x](#)) arrive as mode `100755`; every other file is `100644`.
- **On Windows**, a path Windows can't create stops the clone with a message naming it.

Commit ids are stable: two people who clone the same project get the same ids, and a later fetch never rewrites what you already have.

Staying up to date

Run `git fetch` (or `git pull`) in the clone. Only new commits are converted and downloaded, so a fetch after a day's work is quick. New branches and tags appear like they would from any remote.

If a branch is being saved to while you fetch, the fetch may ask you to retry shortly — it won't build a snapshot of a tree that's still moving.

When a branch's newest commit can't be served yet. A commit enters the git view once it has been *derived* (the Desktop App does this after each commit, or run `swarmfile git index`); a provisional commit (derived but not yet confirmed) is served like any other. A head that has never been derived is derived and claimed on the spot when the credential running git may write — a signed-in app or a key heals the branch, and the command then works (any clone, fetch, push or `ls-remote` heals it; a clone with `swarmfile.confirm=false` has opted out). When it can't (an anonymous or read-only clone, an opted-out clone, or a chain longer than the helper derives in one command), or when the claim is disputed, the branch is left out of the branch listing with that reason printed, rather than serving a stale pointer. Clones and fetches of the project's other branches still work, and asking for that branch by name fails. The lasting fix for those is `swarmfile git index` from a machine with the Desktop App; it derives the selected drive's branch. A disputed claim means two derivations of the same commit disagree. Swarmfile then derives the commit again on the server from its own records and keeps whichever claim matches, so a dispute usually resolves on its own within a minute or two and the branch serves again. The steps below are for the rare dispute where the server's derivation matches neither claim. Running `swarmfile git index` shows which hash your machine reproduces without submitting anything, and a project owner or admin can read both recorded values with `swarmfile git disputes <commit number>` or in the **Commit disputes** section of the project's Commits page on the web. They then reset the commit with `swarmfile git clear-claim <commit number>` or the **Clear** button on that page, using the number the error names. The dispute record is kept after the clear. Swarmfile then derives the commit again on its own (running `swarmfile git index` does the same, sooner), except when Swarmfile's own derivation was part of the dispute: then the commit waits for `swarmfile git index` from a member's machine. Owners and admins are notified when a dispute is settled. A settlement that left the commit provisional can be overridden with the same clear; a settlement that confirmed the commit is final. Each settlement and each clear is recorded, with the values involved, in the **Recent changes** panel of the project's Permissions page (with the whole project selected).

Pushing

`git push` sends commits from the clone back to the project. Each git commit becomes one Swarmfile commit, and Swarmfile keeps the commit's exact author, committer, dates, message and any signature, so its id doesn't change: anyone who clones or fetches afterwards gets the same commit ids you have. Pushing works on managed and unencrypted projects; it isn't available on end-to-end-encrypted ones.

```
git push origin main           # fast-forward an existing branch
git push origin my-feature     # a new branch is created on Swarmfile
git push origin :my-feature    # archive a branch
git push origin v1.0           # a lightweight tag becomes a Swarmfile tag
git push -f origin v1.0        # move the tag to the commit you tagged
git push origin :refs/tags/v1.0 # delete the tag
git push -v origin main        # also print each commit's planned changes
git push --dry-run origin main  # run every check, write nothing
```

A successful push prints a link to the project's commits in the web app (or to the merge request, below).

What a push accepts:

- **Fast-forwards only.** Your commits must build on the branch as Swarmfile has it, one first parent after another. A force push of a branch (`--force`, `+refspec`) is refused, and so is a `git merge` of the remote into a diverged local branch — a merge commit must sit *on top of* the remote head, so fetch and `git pull --rebase` instead. If someone saved to the branch since you last fetched, git reports `fetch first: run git pull --rebase`, then push again. Saves on the drive that haven't been rolled into a commit count too, since a fetch turns them into a snapshot commit. If you need the branch's *tree* to be yours anyway, without rewriting history, run `swarmfile git force-forward` in the clone: it pushes a forward commit whose tree is your tip's (your commits keep riding the auto-archived `merged/<sha>` side branch, so nothing is lost and teammates fast-forward) and prints the `git fetch <remote> && git reset --hard <remote>/<branch>` that re-aligns your clone. A rewritten **root** commit (an amended initial commit) cannot land this way — the replaced history shares no commit with Swarmfile, so the push refuses naming the unrelated root; push that as a new branch instead.
- **New branches.** Pushing a branch Swarmfile doesn't have creates it, starting from the newest commit in your history that Swarmfile already holds.
- **Merge commits** with two parents. If the history you merged isn't on Swarmfile yet, the push sends it first, to a branch named `merged/<first 12 characters of its newest commit>` that starts where that history left the main line (an existing one is reused). The push lists each branch it creates. Once the merge itself lands, the push archives the branch; its commits stay reachable through the merge, so the history doesn't change. If those commits aren't confirmed yet (below), the branch stays until a teammate's fetch has confirmed them, and your next fetch or push archives it then. A `merged/*` branch someone has added commits to is left alone. This works for merges nested inside merged history too. History that shares no commit with the project at all (a merge of an unrelated repository) is refused, naming the commit. Merges of three or more branches (octopus merges) aren't supported.
- **Protected branches** don't move. Your commits go to a branch named `push/<branch>/<first 12 characters of your commit>`, started at the protected branch's head, and the push opens a merge request into the protected branch, or reuses the one already open from that branch. Git reports the ref as not updated, with opened merge request `#<n>` (`push/<branch>/<commit>`); the link takes you to it. The same rules apply to `merged/*` branches: if one is protected, the push is refused.
- **Root creates need root authority.** On a protected project, a commit that adds a new name directly at the project root — or moves one there — is refused unless you're an org owner, the project's creator, or hold a project-wide `write` (or higher) grant; a folder-scoped grant doesn't qualify, and overwriting an existing root file doesn't need it. The push names the remedy: move the files inside a folder and push again, or ask an owner for project-wide write access. See [Permissions](#).
- **Deleting a branch** (`git push origin :<branch>`) archives it, and it can be restored from the web app. `main` and protected branches can't be deleted this way.
- **Tags.** A lightweight tag on a commit Swarmfile holds becomes a Swarmfile tag on that commit. To move an existing tag, retag and push with `--force` (`git push -f origin <tag>`); the commit must already be on Swarmfile. `git push origin :refs/tags/<tag>` deletes a tag. A tag that a [release](#) was published from can be neither moved nor deleted. Annotated tags are refused (Swarmfile tags are lightweight).
- **Large commits.** A commit that changes more than 5,000 paths is sent in parts and applied as one commit, up to 2,000,000 paths. The branch is locked while it applies. Above 100,000 paths Swarmfile applies the commit in stages after the upload, and the push shows its progress. Anyone reading the branch during that time can see the files that have been applied so far. If the push is interrupted after the parts were sent, push again: it picks up where it left off. If someone else's change gets onto the branch while a staged commit is applying, Swarmfile stops the commit. The push then reports how far it got, and the changes already applied stay on the branch as uncommitted changes for you to review. To undo them, use **Roll back** on the notice the web app's Commits page shows for that branch: every file the push touched goes back to how it was before the push, as one new commit you can revert. The person who pushed and project owners can roll back, and the rollback is refused if any of those files has changed again since the push stopped.
- **A commit already on another branch** (for example a fast-forward of `main` to your feature branch's tip) is recorded again on the target branch with the same commit id, so both branches show the same history.

If your history diverged. Two shapes are the same problem: the branch moved while you worked, and a `git merge` of the remote into your line cannot be pushed — a push must continue the remote head along its first parents:

```

A---B---C           the branch on Swarmfile (C is the head)
  \
   X---M           your clone: X, then M = merge(X, C)    <- refused
```

`git pull --rebase` is the supported way back: your work replays on top of `C` as `A---B---C---X'`. When you need your merged *tree* as it stands, `swarmfile git force-forward` lands it as a new commit `F` on top of `C` instead — nothing is rewritten, `C` and `X` stay reachable, and everyone else fast-forwards:

```

A---B---C-----F      F's tree = M's tree
  \                   /
   X-----M
```

A clone can also make plain `git push --force` (and `--force-with-lease`) take the same path: `git config swarmfile.force-forward true`, off by default. Because git's own bookkeeping then points your remote-tracking ref at the tip you sent rather than the commit that landed, follow the note the helper prints — `git fetch`, then reset your branch to the fetched one — before pushing again.

Push an existing repo into a new project. A repository that has never been on Swarmfile goes into a project where the branch you're pushing to (`main` in this example) has no commits and no files yet, with its whole history:

```
git remote add origin swarmfile://<org>/<project>
git fetch origin                # configures git-LFS (and its credential) for Swarmfile
git lfs push --all origin main  # if the repository uses git-LFS
git push origin main
```

The first commit lands with no parent, merges bring their merged history along as described above, and a fresh clone afterwards has the same commit ids. The push sends consecutive commits together, up to 100 commits or 5,000 changed paths per request, within the plan's request budget. Uploading file content and indexing each commit take most of the time. When the budget does run out, the push waits for it and carries on. A `main` that already has files but no commit can't take a root commit: fetch first, and the push builds on the snapshot.

What a pushed commit may contain. Swarmfile accepts a commit only if it is exactly what a clone of the result would produce, so `git fetch` never rewrites it. The push checks every commit, and asks the hub to check the first one, before uploading anything, and names the file and the fix when one doesn't fit:

- **No submodules.** A submodule (mode `160000`) is refused. An executable file (`100755`) is fine: it lands as an executable Swarmfile file, and a commit that only flips the bit (`git update-index --chmod=+x|-x <file>`) is a change like any other.
- **Large files must be LFS pointers.** A file of 1 MiB or more (the project's frozen threshold) must be committed as a git-LFS pointer, and its content must already be on Swarmfile. The clone is configured to send LFS content to Swarmfile with the same credential the clone uses (your signed-in Desktop App session, a stored `sf_key_` key, or `SWARMFILE_API_KEY`; nothing to set up), so git-LFS's own pre-push step normally takes care of that; if a push says an LFS object is missing, run `git lfs push origin <branch>` and push again. A file under the threshold must be committed as a regular file, not as a pointer.
- **Each LFS file is covered in its folder's `.gitattributes`**, either by a `git lfs track` pattern in that folder's own `.gitattributes` (for example `*.psd filter=lfs diff=lfs merge=lfs -text`), or by the line a clone writes for a file the patterns don't cover: `/<name> filter=lfs diff=lfs merge=lfs -text`, one per file, sorted by name, at the end of the file after the line `# swarmfile: generated LFS attributes - view v1`. Only patterns in the file's own folder count, not ones in parent folders, and a pattern with a `/` in it covers nothing. That matters for `git lfs track`: run at the repository root as `git lfs track "media/*.psd"`, it writes a pattern with a `/` into the root `.gitattributes`, which covers nothing here. Put the pattern in the folder's own file instead, for example `printf '*.psd filter=lfs diff=lfs merge=lfs -text\n' >> media/.gitattributes`. A `lockable` attribute on the line (what `git lfs track --lockable` adds) is fine. When a folder's lines don't match, the push prints the lines it expects.

What happens to the files on the drive:

- Changed, new and deleted files become the same changes on the drive. A file renamed without changes keeps its Swarmfile history, and so does a whole folder renamed or moved without changes: it moves as one folder.
- Folders you add are created, and folders your commit empties are removed. Empty folders that exist only on the drive (git can't show them) are left alone.
- A push of several commits sends them in batches, and each commit is applied as a separate Swarmfile commit, in order. If one is refused, the commits before it have already landed, and the branch stays at the last of them. Fix the refused commit and push again to continue from there.
- Commits you push are *provisional* until a second, independent derivation confirms them. Swarmfile's server derives every pushed commit a couple of minutes after the push and confirms it when it reproduces the same commit (a teammate's fetch or a project owner with self-confirmation turned on also confirms). Provisional commits are cloned and fetched like any other. While one is provisional, its row on the project's **Commits** page shows no commit hash — hovering the dash explains why — and the hash appears on the page's next refresh once confirmation lands (the page keeps re-checking for a couple of minutes after a push).

How a pushed commit gets confirmed

A commit is confirmed when a second person works out its file tree independently and gets the same result as the person who made it. You don't have to do anything for that to happen:

- **A teammate's `git fetch` or `git clone` confirms what it brings in.** After the fetch, the helper recomputes each new provisional commit that someone else made, from the branch's history as of that commit, without reusing their results, and sends the result to Swarmfile. Parents go before their children, so a whole chain of commits is confirmed in one fetch. One fetch confirms up to 200 commits, oldest first, and the next fetch continues. A fetch never confirms your own commits, a shallow clone only confirms the commits it fetched, and if confirming fails (for example because of a rate limit), the fetch still succeeds and the next one tries again.
- **swarmfile `git index` confirms a whole branch at once** on a machine with the Desktop App. With no argument, it confirms the branch head and every unconfirmed commit behind it, oldest first, back to the nearest commit that is already confirmed (or one you made yourself). It handles up to 1,000 commits in one run and lists what happened to each commit. `swarmfile git index <seq>` handles just that one commit.

If the two results don't match, the commit can't be confirmed. Swarmfile records both results for the project owner to review. A fetch prints a warning and still succeeds. `swarmfile git index` names the commit and exits with code `2`. Commits that come after it in the same run aren't confirmed.

To keep a clone from confirming commits (a read-only mirror, say), turn it off in that clone:

```
git config swarmfile.confirm false
```

or clone with `git clone -c swarmfile.confirm=false swarmfile://<org>/<project>`.

The same setting also turns off the on-demand derivation of a branch head that has never been derived: a git command in that clone then leaves the branch out of the listing with the `swarmfile git index` guidance instead of claiming it.

When a clone fails

The helper always says why: its message prints above git's `fatal: remote helper 'swarmfile' aborted session` line, and the same text is written to `git-remote-last-error.log` in the Swarmfile cache directory (`~/.cache/swarmfile` on macOS and Linux, `%LOCALAPPDATA%\Swarmfile` on Windows) — include it if you contact support.

The common causes: git access isn't turned on for the project ([Turning on git access](#)); the signed-in account isn't a member of the organization; or the project name in the URL doesn't resolve (use `swarmfile git url <project>` to print a URL that does).

If a clone gets into a bad state

A clone is a local, derived view: every commit id is recomputed from the project's branch state, so a clone you have wrecked — bad refs, a half-finished rewrite, a corrupted object store — can always be replaced. Only work that exists nowhere else is worth rescuing first.

- **Save unpushed work:** `git bundle create rescue.bundle --all` (one file you can move anywhere), or `git format-patch origin/<branch>..HEAD -o ../patches`. To land it on Swarmfile instead, push it to a throwaway branch — `git push origin HEAD:rescue/<name>` creates a new branch forked at the newest commit Swarmfile already holds.
- **Re-clone:** `rm -rf <clone> && git clone swarmfile://<org>/<project>` gives a pristine view with identical commit ids. Generated clone config (git-LFS wiring, `swarmfile.confirm`) is recreated.
- **Repair in place:** `git fetch origin --prune`, then `git checkout -B <branch> origin/<branch>` (or `git reset --hard origin/<branch>`) and `git clean -fdx`; `git fetch --unshallow` if the clone is shallow; `git replace -d <ref>` followed by `git reflog expire --expire=now --all && git gc --prune=now` drops local rewrites.

If a bad state reached Swarmfile itself (junk branches, a moved tag, a bad commit), a local reset won't undo it: archive the junk branch (`git push origin :<branch>`, restorable from the web app), push the tag back (`git push -f origin <tag>`, unless a release used it), commit a revert forward, or ask an owner to clear a stuck commit claim with `swarmfile git clear-claim <seq>`.

What's not supported

Force pushes to a branch, annotated tags	Refused by default. Swarmfile history is append-only; push new commits instead. To land a force push's end state — the branch's tree is yours — as a forward commit, run <code>swarmfile git force-forward</code> in the clone, or set <code>git config swarmfile.force-forward true</code> so plain force pushes are translated (see Pushing). Tags can be moved with <code>--force</code> and deleted, unless a release was published from them. See Pushing for what a push accepts.
Pushing to an end-to-end-encrypted project	Not supported.
Shallow clones	<code>--depth</code> , <code>--deepen</code> and <code>--unshallow</code> work once the commits at the cut have been indexed (<code>swarmfile git index</code>). A depth fetch also needs those boundary commit ids <i>confirmed</i> ; on a project only one person derives, set <code>SWARMFILE_GIT_ACCEPT_PROVISIONAL=1</code> where you run the clone (or let a second person confirm them: a teammate's fetch, or <code>swarmfile git index</code> — see How a pushed commit gets confirmed). <code>--shallow-since</code> and <code>--shallow-exclude</code> aren't supported.
Partial clones	<code>--filter</code> isn't supported. Large files are already LFS pointers, which gives most of the same benefit.
Fetching a commit by id	Only branches and tags can be fetched, not an arbitrary commit id.
HTTPS clone URLs	Not available; clone through the <code>swarmfile://</code> helper.
End-to-end-encrypted projects	Not supported. Mount the project, or use <code>swarmfile materialize</code> to write a ref to a plain directory.
Branch or tag names git can't represent	Refused with a message naming them.

Two kinds of problem tell you how to fix them: a large file that has no recorded id (enable records missing ids automatically, or run `swarmfile git backfill-oids` yourself), and a clone that predates the current git view may need to be re-cloned.

See also

- [Clone and push a Swarmfile project with git](#) — the announcement post, with the end-to-end walkthrough.
- [Git and Swarmfile](#) — every way git and Swarmfile meet, side by side.
- [Coming from Git](#) — the command cheat sheet for day-to-day work on the drive.
- [Runner \(Headless CI\)](#) — CI that runs on a Swarmfile branch without a clone at all.

Git-LFS

Swarmfile can act as a **git-LFS server**. Keep your source in the git host you already use — GitHub, GitLab, or self-hosted — and route only the large binary assets to Swarmfile, where they get real storage, deduplication, encryption (with a documented carve-out for some upload paths — [see below](#)), and quota instead of bloating your git remote. LFS objects count as storage on your plan — deduplicated like drive content — and downloads are free.

This is the lowest-friction way to start: nothing about your VCS or your team's workflow changes. You keep `git push`; `git lfs push / pull / checkout` move the big files through Swarmfile.

The fuller destination is making Swarmfile the home for the *whole* project — native version control on a drive you mount, with no clone at all (see [Coming from Git](#) and [Version Control](#)). The LFS bridge is where you start; you can move the rest at your own pace.

Which projects it works on

git-LFS works on **managed-key** and **unencrypted** projects. It is **not** available on end-to-end-encrypted projects: a stock `git-lfs` client can't hold a per-user encryption key, and the hub structurally can't read end-to-end content, so there is no one to encrypt or decrypt for. An LFS request against an E2E project is refused up front with a `422` and a clear message rather than failing silently per object. See [End-to-End Encryption](#) for what that tier trades off.

Encryption at rest depends on which upload path the object takes, not only on the project's tier — git-LFS is the one place where a managed project can hold plaintext, by design (LFS is an add-on surface, and the direct paths trade encryption for speed on multi-GB artifacts). The paths:

Upload path	At rest on a managed project
Stock <code>git-lfs push</code> (the <code>basic</code> transfer)	Not encrypted — objects are stored directly, up to about 5 GiB each
Desktop app's built-in agent, ordinary sizes (< 64 MiB)	Encrypted — client-side, before upload
Desktop app's built-in agent, ≥ 64 MiB	Not encrypted — uploaded as a single large object; use the standalone agent (or the drive) when encryption is required
Standalone <code>swarmfile-lfs</code> agent (any size)	Encrypted — the agent sends plaintext blocks and the hub encrypts them on write

An **unencrypted** project — a public project on any plan, including every Free-plan project — is plaintext on every path by definition, and end-to-end projects don't support git-LFS at all (above). Objects pushed through the chunked block path also deduplicate against the same content on your mounted drive; the direct paths trade that cross-surface dedup for speed.

If a binary must be encrypted at rest, keep it on the mounted drive (managed or E2E) rather than routing it through LFS.

Importing an existing repository that uses git-LFS is not supported yet: `swarmfile git import` refuses a default branch that carries LFS pointers and skips a side branch that does (see [Known Limitations](#)). Keep the repo on its current host and use the LFS redirect below instead.

Setup

1. Generate a git-LFS credential

In the dashboard, open your project → the **git-LFS** tab → **Generate git-LFS credential**. This mints a project-scoped API key (an `sf_key_...` string). Treat it like a password; it grants access to this one project's files. It counts against your plan's headless-key allowance (Free 1; Starter 2 per seat, minimum 5; Pro 5 per seat, minimum 10) — if the mint is refused, revoke an unused key or add a key pack under Settings → Billing. One key per project is the natural pattern: reuse it for that project's LFS, CI, and scripts rather than minting a fresh one each time.

2. Commit `.lfsconfig`

The panel shows an `.lfsconfig` block. Commit it to the root of your repo so every clone points its LFS traffic at Swarmfile:

```
[lfs]
url = https://hub.swarmfile.com/orgs/<org>/lfs/<project>
```

The URL encodes your org and project, so there's no coupling to your git remote — the git remote stays exactly where it is.

Keep it to just the `url`: `git-lfs` only accepts a small safe allowlist of keys in a repo-committed `.lfsconfig` and prints a `warning: These unsafe '.lfsconfig' keys were ignored` line for anything else — per-machine settings (the large-download retry window, credentials) go in `git config` in the next step instead.

3. Store the credential

`git-lfs` authenticates over **HTTP Basic**. Run the setup once per clone: register git-lfs's filters (without this, `git add` commits the raw bytes instead of an LFS pointer), store the `sf_key_...` in your git credential store as the **password** (the username can be anything), tell git-lfs to use basic auth directly so it skips a 401-probe round trip on every call, and widen the large-download retry window (see [Large files](#)):

```
git lfs install

# Store the credential (username arbitrary, password = your sf_key_...)
git credential approve <<EOF
protocol=https
host=hub.swarmfile.com
username=swarmfile
password=sf_key_...
EOF

# Use basic auth for this endpoint without the 401 probe
git config lfs.https://hub.swarmfile.com/orgs/<org>/lfs/<project>.access basic

# Per-machine: how many times a large download retries while it's prepared.
# (.lfsconfig can't carry this – git-lfs ignores non-allowlisted keys there.)
git config lfs.transfer.maxretries 240
```

The "Generate git-LFS credential" panel emits this snippet with your real org, project, and key filled in — copy it from there; it has a **Windows (PowerShell)** tab for machines where bash is not available. A `~/.netrc` entry (`machine hub.swarmfile.com login swarmfile password sf_key_...`) works too if you prefer.

On Windows, run the same setup from PowerShell instead — the credential step uses a here-string, since `<<EOF` is bash-only:

```
git lfs install

@'
protocol=https
host=hub.swarmfile.com
username=swarmfile
password=sf_key_...
'@ | git credential approve

git config lfs.https://hub.swarmfile.com/orgs/<org>/lfs/<project>.access basic
git config lfs.transfer.maxretries 240
```

4. Track and push

From here it's stock git-LFS — nothing Swarmfile-specific:

```
git lfs track "*.psd" "*.exr" "*.uasset" # writes .gitattributes
git add .gitattributes
git add my-big-file.psd
git commit -m "Add art source"
git push # source objects to your git host
git lfs push # big files to Swarmfile
```

Teammates who clone the repo get the committed `.lfsconfig`, add their own credential once (step 3), and `git lfs pull` fetches the big files from Swarmfile.

Or let the GitHub App do it

Everything above works with any git host and any git-lfs client. If your remote is GitHub, the project's git-LFS tab can do part of it for you (the git-LFS tab installs the Swarmfile GitHub App and opens the setup pull request; the manual steps above always work). If you belong to more than one organization, the install page asks which one to link the installation to — one installation maps to one Swarmfile organization:

- **Install the GitHub App** — click the button in the tab, pick the repositories to let Swarmfile see, and come back. One installation per Swarmfile organization.
- **Route a repository** — in the same tab, choose the repository and click **Route this repo to <project>**. That links it to the project (each repo gets its own `.lfsconfig`, because the URL embeds the project id).
- **Open setup PR** — Swarmfile opens a small pull request that adds `.lfsconfig` to the repository. Merge it (or close it and use the manual steps above — nothing is forced).
- **Mint git-LFS credential** — the tab mints a project-scoped `sf_key_...` for you and shows the same per-clone commands as step 3, prefilled.
- **Filter or bulk-link** — an installation with many repositories gets a filter box and a **Link all shown** action, so you don't link them one at a time.
- **Re-run reflect** — if a push's automatic reflect misses (a very large push, or an object that was still being verified), the repository row's **Re-run reflect** runs it again from the branch head.

After that, `git lfs push/pull` work exactly as described above, and **agent-pushed objects show up on the mounted drive automatically** — a push webhook reflects the pushed branch's objects onto the project's matching Swarmfile branch, so you don't need to run `swarmfile-lfs reflect` by hand. On the machine that pushed, the file appears on the drive within moments as an optimistic preview; the hub's reflect then confirms it, or withdraws the preview if it refuses the mapping (a path conflict, or a folder you can't write to). Teammates see the file once the hub reflect lands, as before. (Objects pushed by a stock `git-lfs` client are pull-able but not reflected — see the note under [Deduplication and the mounted drive](#).) The App is provisioning-only: it never sees or proxies your LFS objects, and the credential stays the `sf_key_...` you minted — uninstalling the App doesn't break clones that already have their credential.

Deduplication and the mounted drive

Swarmfile stores content once. If two people push the same bytes — or the same object already exists — the second upload is skipped, exactly as LFS dedup is meant to work. One exception: when the bytes so far exist only as a file saved on the drive, the first push uploads them anyway, and Swarmfile checks the drive copy in the background; once it's confirmed, later pushes of the same bytes are skipped.

It goes one step further: a file saved to the **mounted Swarmfile drive** through the desktop app is `git lfs pull`-able by its object id, with **zero duplicate storage**. (So is a file added through the web dashboard's uploader. A file saved before object ids were recorded can be given one with `swarmfile git backfill-oids`.) The LFS object is the native content — the same bytes the drive already holds, addressed two ways. So a file your team produced on the drive can be pulled by a git-LFS client that only knows its `oid`, without storing it twice.

That holds for any version the project still keeps: a drive version that has aged out of your history retention is gone for LFS too, and a pull of it reports the object as missing. On the default **Keep every saved version** projects a drive version doesn't age out, so its LFS object stays pull-able while the version is retained. Pushing or re-pushing an object with `git lfs push` makes it an LFS object in its own right, kept independently of drive retention. It doesn't apply to end-to-end encrypted projects (no LFS at all), nor to content encrypted with a static key you supply yourself (`SWARMFILE_ENCRYPTION_KEY`): in both cases the hub holds no key, so Swarmfile never sends it a file's SHA-256 object id.

The reverse direction works too, with one manual step. An object pushed with `git lfs push` can be **projected onto the mounted drive** as a first-class file at its repo path — reusing the blocks the push already stored (again, zero duplicate storage). Because the LFS protocol only ever sends an object's *id*, never its path, the file's path comes from your git working tree, so you run one command after pushing:

```
git lfs push origin main
swarmfile-lfs reflect          # projects the just-pushed objects onto the drive
```

`reflect` reads `git lfs ls-files`, sends the oid→path map to Swarmfile, and the hub creates the drive entries. It mirrors your current git branch onto a same-named Swarmfile branch (auto-created; override with `--branch`), and it **never overwrites** a file already on the drive — a path already occupied by different content is reported as a conflict and left alone. Re-running it is safe (already-projected files report as unchanged). (If you run the desktop app rather than the standalone binary, `swarmfile-engine reflect` is the identical command.) To make it automatic, wire it into a `git push` alias:

```
git config --global alias.pushr '!git push "$@" && swarmfile-lfs reflect'
# then: git pushr origin main
```

After reflection, that path has a real drive entry — so it's reachable through *both* surfaces (drive + LFS), participates in mount worksharing locks, and appears for teammates on the drive.

Reflection needs the transfer agent. `reflect` maps an object onto the drive only when Swarmfile stored it as content it can address — i.e. the object was pushed through the transfer agent (above) or the desktop app. A plain **stock git-lfs push** takes the direct-upload path: the object is stored and `git lfs pull`-able, but it isn't a drive manifest, so `reflect` reports it as `missing`. If a file must appear on the mounted drive, install the agent (see [Pushing large files](#)) and push through it, or write the file to the drive directly. The GitHub App's automatic reflect has the same dependency.

Reflection is **additive**: it creates or leaves-alone, and never deletes. Removing or renaming a file in git won't remove or move its drive entry — do that on the drive (or via the dashboard). It enforces the same **write** permission a drive write does: on a project with per-folder access rules, reflecting a file into a folder you can't write to is refused (reported as `denied`) and skipped.

A very large push reflects a **bounded number of objects per run** (a whole-tree scan is capped for cost). The rest land on the next push — or immediately, via another `swarmfile-lfs reflect` (or the tab's **Re-run reflect**). The cap is noted in the row's status when it applies.

Large files

Small and mid-sized objects download immediately. A **very large** object (roughly a gigabyte or more) is prepared on its **first** download: Swarmfile assembles it into a single downloadable copy in the background, and until that finishes `git lfs pull` waits and retries automatically — you'll see it pause, then complete. This is why step 3 sets `git config lfs.transfer.maxretries 240`: it gives the client enough patience to wait through preparation. Once prepared, that copy is reused, so subsequent pulls of the same object are immediate. Under the hood, an attempt while preparation is still running gets `429 Too Many Requests` with a `Retry-After` interval — that pause is git-lfs honoring the retry — and a preparation that fails for good gets `502` naming the reason rather than looping the client on 429.

If a pull of a very large object ever gives up with a "being prepared, retry" message, just run `git lfs pull` again — preparation continues in the background and the retry will pick up the finished copy. A `502` is different: preparation failed permanently and the message says why, so retrying won't change the outcome until the underlying cause clears. (Raising `lfs.transfer.maxretries` further with `git config` widens the wait window if you routinely pull multi-gigabyte objects. It has to be `git config`, not `.lfsconfig` — git-lfs ignores that key in a committed config file.)

Pushing large files

A plain `git lfs push` sends each object in a single request, supporting objects up to about 5 GiB. To push bigger files — or to get chunked, resumable uploads — install the **Swarmfile transfer agent**: a small (~2 MB), standalone, open-source binary (`swarmfile-lfs`) that speaks git-lfs's documented [custom-transfer protocol](#). No desktop app required — but if you already run the desktop app, it's already installed (every desktop installer bundles it on `PATH`), so skip the download and go straight to `swarmfile-lfs install`.

```
# 1. If the desktop app is installed, the agent is already on PATH — skip to
# step 2. Standalone install:
cargo install swarmfile-lfs-transfer          # any Rust toolchain
curl -fsSL https://get.swarmfile.com/lfs | sh  # macOS / Linux
# Windows (PowerShell): irm https://get.swarmfile.com/lfs.ps1 | iex
# ...or download the binary for your platform from the public release:
# https://swarmfile.com/public/swarmfile/swarmfile-lfs
# (Homebrew is not available yet; use cargo or the installer above.)

# 2. Wire git-lfs to it, once per machine:
swarmfile-lfs install
```

Those standalone channels ship separately from the app, so the dashboard's install command is the authoritative one for your account. If you already run the desktop app, `swarmfile-lfs` is installed on `PATH` with it, and the app also bundles an equivalent built-in agent (`swarmfile-engine-lfs-transfer`; its encryption path differs slightly — see below).

`swarmfile-lfs install` just writes the git-lfs custom-transfer adapter config for you (add `--local` to scope it to the current repo instead of your global config; `swarmfile-lfs uninstall` undoes it in the same scope). With it in place, `git lfs push` transparently uploads large files as small content-addressed blocks — encrypted at rest on a managed project (the built-in agent seals blocks client-side; the standalone agent's blocks are encrypted by the hub on write) — up to **256 GiB per object** (and your project quota), deduped against the same content on your Swarmfile drive. Nothing else changes: `git lfs push/pull` work exactly as before, the agent only kicks in for uploads, and it authenticates with the same credential you stored in step 3 (so make sure that's done on each machine that pushes large files).

Without the agent, stock `git lfs push` still works up to about 5 GiB per object, and any file is always available by writing it to the mounted drive instead. If you already run the **Swarmfile desktop app**, `swarmfile-lfs` is already on `PATH` — `swarmfile-lfs install` is all you need. The engine also exposes an equivalent built-in agent as `swarmfile-engine-lfs-transfer`, if you'd rather point git-lfs at that (`git config lfs.customtransfer.swarmfile-chunked.path swarmfile-engine / .args lfs-transfer`). (The built-in and standalone agents differ in which upload path they use — the table above shows which paths are encrypted at rest.)

An interrupted large push is safe to re-run. Just run `git lfs push` again: content-addressed blocks the agent already uploaded are deduplicated, and the desktop app's built-in agent additionally resumes a ≥ 64 MiB multipart upload from the parts that already landed instead of starting over. Nothing is left behind either way — an upload nobody resumes is reaped server-side after 7 days on Swarmfile-managed storage.

What's supported

- The **LFS Batch API** (`upload` and `download`) and the **basic transfer** adapter — the operations `git lfs push`, `pull`, `checkout`, and `fetch` use.
- Verify**, so a truncated upload is caught rather than silently stored.
- Locking** (`git lfs lock` / `git lfs locks` / `git lfs unlock`), for every tracked path — including files that live **only** in your git host and have never been written to the Swarmfile drive. When a path *does* also exist on the drive, its LFS lock and its desktop/mount worksharing lock are the **same** server-side lock, so locking via `git-lfs` and locking on the drive mutually exclude, across machines and across both surfaces. The path resolves on the branch `git-lfs` reports (its `ref` — the checked-out branch) whenever Swarmfile has that branch; otherwise it falls back to the project's default branch, so an LFS-only repo whose git feature branches have no Swarmfile counterpart keeps working as before. A path the resolved branch's tree doesn't have — a file that exists only in git, or one missing from that branch — takes the branch-less advisory lock below. Lock listing and push verification scope entry locks the same way — a drive lock taken on a feature branch doesn't block a push to `main` — while the branch-less advisory lock applies to every branch. The advisory lock is what other `git-lfs` clients see; unlocking is owner-gated (use `--force` to break another user's lock). Locks hold for 30 days and nothing renews a git-side lock, so a lapsed one releases itself rather than staying stuck — take it again if the work outlives that. See [Working with Files](#).

Performance

Upload throughput is bounded by *your* connection to Swarmfile, so no headline number is meaningful. Two properties hold regardless of uplink:

- The chunked agent is faster for large objects**, because it sends many ~ 1 MiB blocks in parallel instead of one large request.
- Re-pushing an unchanged object transfers nothing** — git-lfs sees the server already has the object and skips it, so a re-push costs one batch round-trip (well under a second) no matter the file's size.

Continuous integration (CI)

Two things make CI painless, and they split by direction:

- Pulling assets needs no agent for almost everything.** `git lfs pull/fetch/checkout` works with the stock git-lfs that's already on every runner; large objects materialize and retry (that's what the `lfs.transfer.maxretries` setting below is for). One boundary: a very large object's first *stock* download may be paced — the hub answers with a "wait, preparing" and the retry setting carries the client through; on a deployment without the prepare queue it is refused outright, and the mount is the read path. Only uploads have a size boundary the transfer agent lifts. A build or test job needs only credentials.
- Pushing objects over ~ 5 GiB needs the agent.** A release or asset-generating job adds one install step.

In both cases, store the project's `sf_key_...` as a CI secret (`SWARMFILE_API_KEY`) and feed it to git through a credential helper. Never bake the key into `.lfsconfig` or a command line.

GitHub Actions

Pull-only job (build/test — no agent):

```

env:
  SWARMFILE_API_KEY: ${ secrets.SWARMFILE_API_KEY }
steps:
  - uses: actions/checkout@v4
    with: { lfs: false }          # skip the auto-smudge; we pull explicitly below
  - name: Configure Swarmfile LFS credentials
    run: |
      git config --global credential.helper \
        '!f() { test "$1" = get && printf "username=lfs\npassword=%s\n" "$SWARMFILE_API_KEY"; }; f'
      git config --global lfs.transfer.maxretries 240 # wait through large-object preparation
  - run: git lfs pull             # any size the hub can assemble; a very large object may ask the client to wait while it prepares

```

Push job (release/asset generation — installs the agent). If you can't use the `get.swarmfile.com` installer (a restricted network, or a policy against piping a download to a shell), download `swarmfile-lfs` from the [downloads page](#) or copy it off a machine with the desktop app, and substitute that path for the `curl` step:

```

env:
  SWARMFILE_API_KEY: ${ secrets.SWARMFILE_API_KEY }
  SWARMFILE_LFS_URL: https://hub.swarmfile.com/orgs/<ORG>/lfs/<PROJECT>
steps:
  - uses: actions/checkout@v4
  - name: Install the Swarmfile transfer agent
    run: |
      curl -fsSL https://get.swarmfile.com/lfs | SWARMFILE_LFS_BINDIR="$HOME/.local/bin" sh
      echo "$HOME/.local/bin" >> "$GITHUB_PATH"
      swarmfile-lfs install # writes the git-lfs custom-transfer config (global)
  - name: Configure Swarmfile LFS credentials
    run: |
      git lfs install # register the clean/smudge filters on this runner
      git config --global credential.helper \
        '!f() { test "$1" = get && printf "username=lfs\npassword=%s\n" "$SWARMFILE_API_KEY"; }; f'
      git config --global "lfs.${SWARMFILE_LFS_URL}.access" basic
      git config --global lfs.transfer.maxretries 240
  - run: git lfs push origin HEAD # large objects (>5 GiB) now go through the agent

```

GitLab CI

```

push-assets:
  image: ubuntu:22.04
  variables:
    SWARMFILE_LFS_URL: https://hub.swarmfile.com/orgs/<ORG>/lfs/<PROJECT>
  before_script:
    - apt-get update && apt-get install -y git git-lfs curl ca-certificates && git lfs install
    - curl -fsSL https://get.swarmfile.com/lfs | SWARMFILE_LFS_BINDIR=/usr/local/bin sh
    - swarmfile-lfs install
    - git config --global credential.helper
      '!f() { test "$1" = get && printf "username=lfs\npassword=%s\n" "$SWARMFILE_API_KEY"; }; f'
    - git config --global "lfs.${SWARMFILE_LFS_URL}.access" basic
    - git config --global lfs.transfer.maxretries 240
  script:
    - git lfs push origin HEAD
  # Set SWARMFILE_API_KEY as a *masked, protected* CI/CD variable — not in this file.

```

The credential helper reads `SWARMFILE_API_KEY` at the moment git invokes it, so the secret never lands on a command line or in the log. For a pull-only job you can drop the agent install and the `access basic` line entirely.

Leaving Swarmfile

Your source lives in your own git host and never touches Swarmfile — so leaving is only ever about pulling your binary objects back out, and that works with **stock git-lfs** and **no Swarmfile software installed at all**. There is no export API to learn and no proprietary format to unwind.

```

# In a clone that points at Swarmfile, pull every version of every object
# referenced by any branch or tag into your local .git/lfs/objects store:
git lfs fetch --all

# Confirm you hold complete, uncorrupted copies (each object hashes to its oid):
git lfs fsck

```

After `git lfs fetch --all`, every object is on your disk. To move them to another LFS server, point the repo at the new endpoint and push:

```
git config -f .lfsconfig lfs.url https://your-new-lfs-server/... # GitHub, GitLab, S3-backed, self-hosted, anything
git lfs push --all <remote>
```

This is deliberately the *standard* git-lfs export path — the custom-transfer agent is only an optional accelerator for large **uploads**; nothing about your objects is locked to Swarmfile software. You can walk away with a single command, which is exactly the point: your binaries are as portable as your git history.

Troubleshooting

A warning: These unsafe '.lfsconfig' keys were ignored line. Something other than `lfs.url` is in the committed `.lfsconfig` — older setups put `lfs.transfer.maxretries` there. git-lfs only accepts a safe allowlist in a repo-committed config file and ignores the rest by design; move per-machine settings to `git config` (step 3 shows the retry one). The warning is about the ignored key, not a failed transfer.

A 401 on every LFS call. Credential helpers disagree about whether the key belongs in the username or the password field; Swarmfile accepts the `sf_key_` in **either**, but the safest placement is the password. Re-run the `git credential approve` step above, and confirm you set `..access basic` so git-lfs sends the header on the first request.

A 422 for every object. The project is end-to-end encrypted, which git-LFS can't support (see above). Use a managed or unencrypted project, or move the source-plus-assets to a mounted drive instead.

git lfs pull fetches nothing / objects are missing. Confirm `.lfsconfig` is committed and its `url` names the right org and project, and that the credential you stored is for that same project.

git lfs unlock fails on a lock someone else holds. Path locks are owner-gated: only the user who took a lock can release it normally. Pass `--force` (or `--id <id> --force`) to break another user's lock, or ask them to unlock it. `git lfs locks` shows the current holder of every live lock.

Coding agents (mounted workspaces)

A Swarmfile project mounts as a real filesystem path. That makes it a good workspace for a coding agent: the agent reads and writes ordinary files, and every change is versioned, lockable, and shared — no `git clone`, no `git worktree` per task, and no full re-install of dependencies for every parallel checkout.

The shape this guide builds:

- one **engine process** per machine (or per agent when you want process isolation — see below),
- one **project API key** so the engine runs headless (no browser sign-in),
- one **mount per agent**, each with a label and (optionally) its own branch,
- one **shared scratch cache** per project per machine, so a second mount reuses the first one's downloaded packages,
- and the **durable create queue**, which is what makes a burst of file creations from an agent safe.

The same primitives appear in the Desktop App — this guide is the scripted, headless form.

What each agent gets, and what it must not be given

Give an agent:

- **The mount path.** It's a normal POSIX/Windows path (`/private/tmp/sf-ex-b/...`, `X:\...`). Everything the agent does happens there; there is no separate "sync" step.
- **The scratch dir** (below) as the place for tool caches.
- **Exit codes** (below) so its automation can tell "you need to do something" (2) from "something failed" (1).
- **One `SWARMFILE_*` environment set** per engine — never two engines sharing a cache directory.

Do **not** hand an agent:

- A second mount point that points at the *same directory* another agent uses — mount points are exclusive to one engine process.
- A personal refresh token copied off a workstation. Use a project API key: it's revocable, it's walled off to one project, and it can't act as the human who minted it.

Platform matrix: how many mounts per engine

Platform	Several mounts in one engine?	Recommended shape for N agents
Windows (WinFsp)	Yes — <code>mounts open</code> takes a drive letter (<code>Y:</code> , <code>Z:</code> , ...), <code>auto</code> for the next free one, or an empty folder (a folder with anything in it is refused)	One engine, N mounts, one API key
Linux (FUSE)	Yes — <code>mounts open</code> takes directory paths (same engine path as Windows; macOS/Linux use a folder, not a drive letter)	One engine, N mounts, one API key
macOS (FUSE-T)	Yes — <code>mounts open</code> takes directory paths; multiple live mounts are supported (up to 18 per Mac — see below)	One engine, N mounts, one API key

An engine can serve every agent's mount: auth is one API key, and the shared scratch cache (below) is per machine, so the dependency cache is still paid for once. Running one engine per agent remains a valid isolation choice (a wedged mount helper then can't affect siblings), but it isn't required on macOS. A Mac supports up to 18 simultaneous mounts; if an `open` fails because that limit is reached, close a mount and try again.

1. Headless engine with a project API key

The commands in this guide use the `swarmfile` CLI, which ships with every platform installer (Linux `.deb`, the Windows MSI, the macOS `.pkg`). A headless box needs only that package — no GUI sign-in — and the CLI resolves the running engine's control socket on its own. See [CLI: swarmfile](#) for the per-platform install locations.

Mint a key as an org admin or owner, from any running engine:

```
swarmfile api-key create agent-node-1          # this mount's project
swarmfile api-key create agent-node-2 --project proj_... # or an explicit project id
```

`--project` takes a **project id**, not a slug (a slug is refused with `project_not_found`, and the error names the id to use). The secret is printed **once** (`sf_key_...`). List and revoke keys with `swarmfile api-key list` and `swarmfile api-key revoke <id>`. A key is walled off to exactly one project, so a compromised agent node can't reach anything else in the org.

Keys count against your plan's **headless-key allowance** — Starter 2 per seat (minimum 5), Pro 5 per seat (minimum 10) — and a key pack adds 10 for \$20/month. One key per machine/fleet is the recommended pattern (revocation radius and its own rate-limit bucket); a `402` with `code: keys_cap` means the org is at its cap, so revoke an unused key or add a pack under Settings → Billing. Agents running through a workstation's existing mount need no key at all.

To give each agent its own engine (process isolation; otherwise one engine can host every agent's mount, as the matrix above recommends), point each at its own cache dir and mount point:

```
export SWARMFILE_API_KEY=sf_key_...
export SWARMFILE_ORG_ID=org_...
export SWARMFILE_PROJECT_ID=proj_...
export SWARMFILE_CACHE_DIR="$HOME/.swarmfile/agent-1" # one per engine
export SWARMFILE_MOUNT_POINT="$HOME/agents/agent-1" # one per engine
export SWARMFILE_MACHINE_ID=agent-node-1
swarmfile-engine
```

With `SWARMFILE_API_KEY` set, the engine skips sign-in entirely — no browser, no refresh token. (If both an API key and a refresh token are set, the key wins and the engine says so in its log.) See the [api-key CLI reference](#) for the full flags, and [Deployment Topologies](#) for putting these on dedicated boxes.

2. One mount per agent (and targeting it)

The engine's first mount is the **boot mount**, id `default`. Additional mounts are full, symmetric mounts — each with its own changelists, metadata, and lock manager:

```
swarmfile mounts list
swarmfile mounts open Y: --label agent-2 --branch feature-x
swarmfile mounts open Z: --label agent-3 --project proj_other
swarmfile mounts status agent-2-id
swarmfile mounts current agent-2-id
```

The examples use Windows drive letters; on macOS and Linux the mount point is a directory path instead (e.g. `swarmfile mounts open /Volumes/agent-2 --label agent-2`). Everything else — labels, ids, `--branch`, `--project`, `status`, `current` — is the same on every platform.

- `--label` is for humans; the id (a long hex string, or `default` for the boot mount) is what scripts use — take it from `mounts list --json`.
- `--branch <name>` pins that mount to a branch. It opens (or reuses) a divergent (`project`, `branch`) scope, so the mount works that branch while every other mount stays on its own — and can be switched later with `mounts branch`. A unique directory root per agent (e.g. `/agents/<name>/...`) is still good hygiene, but not required: the engine creates missing ancestor directories on a directory create, so a fresh branch is safe to write into anywhere.
- `--project <id>` opens a mount on a different project in the same org.
- `--sparse <glob>` (repeatable) opens a **sparse** mount: only the matching paths are visible (and writable) on that drive, so an agent working one package never sees the rest of the repo. `--exclude <glob>` is sugar for the `!` negations ("everything except"); an excluded path has no listing entry, refuses an open and a write-create, and the mount's offline (`hydrate`) set matches. Scope is per mount, so branches mount differently.
- `--exclusive` opts that mount into **exclusive-write mode**: while it has a file open for writing, another mount opened with the same flag on the same project+branch has its write-open refused (the app sees `EAGAIN`, or "file in use" on Windows) — a rename, delete, or atomic-save replace of that file is refused too — instead of both editing the file and the later save becoming a conflict to resolve. Off by default, fixed at open, and shown as a lock badge in the drive list. Open **every** agent mount with it when several agents share one tree: two mounts of one engine are a single identity to the hub, so ordinary locks do not separate them. If a successful open warns that the option wasn't applied, update the engine.
- Every command that isn't under `mounts` targets the current mount unless you pass the global flag:

```
swarmfile --mount agent-2-id status
swarmfile --mount agent-2-id commit -m "agent 2 results"
```

Close a mount with `swarmfile mounts close <id>`. It refuses while a changelist is open (`--force` abandons it) and refuses the boot mount outright (`exit 2`, `is_default_mount`) — closing that one means shutting down the engine.

3. Shared dependency cache

Every mount of a project on a machine can share one **scratch directory** for tool caches (a pnpm store, a cargo registry, a pip wheel cache). Declare what `run` should redirect, in a committed file at the project root:

```
# .swarmfile/cache.yml
caches:
  - name: pnpm
    env: PNPM_HOME
    subdir: pnpm
  - name: cargo-registry
    env: CARGO_HOME
    subdir: cargo
```

`swarmfile scratch init` writes a commented starter file. Then `run` installs and builds through `run`, so the tool sees the shared cache:

```
swarmfile mounts scratch-dir default # the shared dir + declared redirects
swarmfile run -- pnpm install # PNPM_HOME -> shared scratch
swarmfile run -- cargo build --release # CARGO_HOME -> shared scratch
```

`run` passes `stdio` and the exit code through, and it never refuses to run your command over a cache problem — with no config, a bad entry, or an unreachable engine it warns and runs the command unmodified.

Maintenance:

```

swarmfile scratch list
swarmfile scratch gc --older-than-days 14 --dry-run
swarmfile scratch clear <name> --force

```

One scratch dir per project per machine, shared regardless of branch. Only redirect **read-mostly, content-addressed caches** — not build output directories, which would contend between concurrent mounts. Nothing in a scratch cache is unrecoverable project data; clearing one only costs a slower next build. Full reference: [scratch](#) and [run](#).

4. Create bursts and the pending-create queue

File creation returns immediately on the drive. On the mounted drive, a create (or `mkdir`, or a symlink on macOS/Linux — Windows refuses symlink creates) returns at once under a locally minted id; a background worker confirms it with the hub (8 in flight by default) and peers on the LAN see the name by gossip before the hub does. Creates that are due together go to the hub together, up to 50 in one request, and a new file's first save travels with its create instead of in requests of its own. This is what makes an agent's scaffolding burst safe and fast: 5,000 new files reach the hub in under half a minute, and a create made while the hub is unreachable simply queues.

What to check, and what an agent will see:

```

swarmfile status          # "new files: N waiting to reach the hub (oldest ...)" while queued
swarmfile --json status | jq .pendingCreates

```

pendingCreates field	Meaning
queued	Creates this machine made that the hub hasn't confirmed. <code>0</code> = drained.
canceled	Deleted locally before the hub heard of them; no hub call needed.
stuck	Queued and failed 6+ times — hub refusing or unreachable for these.
oldestAgeSecs	Age of the oldest queued create. Growing with <code>queued</code> unchanged means nothing is draining.
peerUnconfirmed	Entries peers gossiped as created that this machine hasn't seen confirmed.
items	The oldest 20 queued creates (path, type, attempts, age).

Rules of the road while something is pending:

- **A name collision on the hub renames, it doesn't overwrite.** If the hub already has a different file at that path when the queued create lands, your copy is kept as `name (conflict-xxxxxxx).ext` and the error ring says so. A **folder** with the same name is merged into the hub's folder (like `mkdir -p`) — files you created inside it land there, and the hub's existing contents appear after a few seconds.
- **Operations on a just-created file wait, then answer 409.** Locking, ACL reads, share-link creation and comment writes on a pending file wait up to 8 seconds for the create to confirm and then proceed. A change (a lock, share, ACL or comment write) sends the file's create at once, even while its first save is still uploading, so it doesn't wait behind that save. If the hub can't be reached in that time they answer 409 with `code: "pending_create"` — retry shortly. For an entry *another* machine created and the hub hasn't confirmed, the refusal is immediate.
- **A create the hub refuses for good is retracted** from peers and reported once in the error ring (`create failed for "..."`). An empty file, a symlink, or a folder is removed. A file you had already written content into is **kept on this computer only** (never uploaded) so the work isn't lost — the message says so. Content in files under a refused folder is discarded, and the count is in the message.
- **Deleting or renaming a just-created file** is passed to other machines immediately. A file a peer created that never reaches the hub is removed after 6 hours with a message.
- **A new file can reach the hub up to a minute after it appears locally.** While a new file's first save is still uploading, its create waits for it (up to 60 seconds) so the two go to the hub as one request. Other machines on the LAN see the file at once by gossip, and the web dashboard's Files view shows "N files arriving from ..." until the hub has them.

Tunables, rarely needed:

Variable	Default	Effect
<code>SWARMFILE_CREATE_DISPATCH_CONCURRENCY</code>	8 (1–64)	Creates sent to the hub at once when they go one by one. Lower it if the hub keeps answering 429; raise it on an unthrottled hub.
<code>SWARMFILE_PENDING_CREATE_WAIT_SECS</code>	8 (0–120)	How long lock/ACL/share/comment writes wait for a pending create before 409 <code>pending_create</code> . <code>0</code> refuses at once.
<code>SWARMFILE_CREATE_BATCH</code>	on	Send due creates in batches of up to 50 per request. <code>0</code> sends each on its own.
<code>SWARMFILE_CREATE_COALESCE</code>	on	Send a new file's first save with its create (one request). <code>0</code> sends the create, then locks and commits separately.
<code>SWARMFILE_CREATE_HOLD_UPLOADING_SECS</code>	60 (0–600)	How long a create waits for its first save while that save is uploading. The upload finishing sends it at once, so this only caps a stall.
<code>SWARMFILE_CREATE_HOLD_SECS</code>	2 (0–10)	How long a create waits for a file that is open for writing but not saved yet. <code>0</code> turns off all waiting.
<code>SWARMFILE_ARRIVING_REPORTS</code>	on	Tell the project how many new files are on their way (at most once every 30 s), for the dashboard's "files arriving" note. <code>0</code> / <code>false</code> / <code>off</code> turns it off.

These are client-side. The hub's own ceilings are per-org and per-user (plan defaults, owner/admin configurable), and a single API key can carry its own request limit (`swarmfile api-key rate-limit`) — the right lever when a fleet of agents shares one key and needs more headroom than one person's window. `swarmfile admin rate-limits usage` reports this minute's usage against them; see [Organizations, Projects & Members](#).

How fast can one user create files?

The per-user ceiling is a number of requests per minute, and it depends on the plan. Defaults:

Plan	Per-user default	Owner/admin can raise it to	Per-org default
Free	100 / min	fixed	300 / min
Starter	300 / min	500 / min	1,000 / min
Pro	500 / min	1,000 / min	1,000 / min
Enterprise	500 / min	5,000 / min	10,000 / min

What counts against it is requests, not bytes, and block uploads are not counted at all:

Kind of new file	What it costs
A new file with content (saved when it is created)	one create that carries its first save: 1 request , or a tenth of one when it goes out in a batch
An empty file or folder	1 request , or a tenth of one in a batch
A file created empty and written later	create, take the write lock, then one commit that also gives the lock back: 3 requests

Batching is paced to the plan's per-minute allowance: on Pro's default ceiling (500 requests a minute), a 5,000-file burst drains in well under a minute — batched creates cost about a tenth of a unit per file, so a full 50-file batch is 5 units even though it is a single HTTP request. Nothing is lost when a limit is reached; the queue simply drains more slowly.

If a fleet needs more, give it its own API key with its own ceiling (`swarmfile api-key rate-limit`) instead of raising everyone's, or ask an owner or admin to raise the org and user limits (`swarmfile admin rate-limits` ; a raise can be time-boxed with `--for 6h`).

5. Parking and switching

An agent's scope can move without losing queued or staged work:

- **Branch switch (hot).** `swarmfile branch switch <name>` moves the boot mount; `swarmfile mounts branch <id> <branch> [--wait]` moves one independent mount. Both refuse (exit 2) while a changelist is open. Creates queued before the switch land on the branch they were made under — the scope is captured when the create is enqueued, not when it drains.
- **Project/org switch.** `swarmfile workspace switch --org <id> [--project <id>]` moves the whole engine. A same-org switch usually takes a hot path; the fallback restarts the engine process (the OS supervisor relaunches it). Queued creates and their corrective renames/deletes stay in the project captured at enqueue.
- **Park.** `swarmfile changelist park` sets staged work aside so a branch or project can be left; `swarmfile changelist parked` lists it (age, files) and `swarmfile changelist resume` restores it. Parked work is a *set* per branch, and it's kept alive on the hub, so it survives the switch. `changelist discard --yes` abandons it.
- **Upstream refresh.** A branch is a fork: it does not see later work on `main`. There is no `main`→branch merge exposed — `swarmfile merge` and `mr` both go *into* the parent — so plan agent branches to be short-lived, or recreate the branch from current `main` when it needs the latest. A merge-down surface is a known gap.
- On a TTY, `workspace switch` prompts before abandoning staged/queued work and offers to park first; `--yes` / `--json` skip the prompt for scripts.

The tray mirrors all of this: the mount switcher lists every mount with its branch and per-row switch — and the tray menu's **Open Another Drive...** opens the same **Open a drive** dialog (its **Prevent another drive from editing a file this one has open** checkbox, under **Advanced**, is the same `--exclusive` opt-in), so a second mount is one menu click away — while a **Parked changes** panel lists, resumes (all at once or one at a time), and discards parked work. The dialog puts the drive at the next free drive letter on Windows and in a folder next to your main drive on macOS and Linux, unless you choose otherwise under **Advanced**. If a mount can't open, the form says why straight away — for a drive letter that's taken it offers the next free one — and the failed mount never becomes the current one.

6. What to hand an agent (a copy-paste preamble)

```
You are working in a Swarmfile mount, not a git checkout.
- Workspace path: <mount path>                (writes sync automatically; there is no commit-required step)
- Tool caches:   run builds via `swarmfile run -- <cmd>` from the mount
- Status:        `swarmfile status` / `swarmfile --json status` (jq .pendingCreates)
- Exit codes:    0 success, 1 failure, 2 a refusal you should act on,
                 3 a long operation detached (`swarmfile op status` follows it)
- If an operation answers 409 with code "pending_create", the file is still
  syncing: wait a moment and retry. Do not delete and recreate it.
- Don't run two engines against the same cache dir or mount point.
```

Runner (Headless CI)

A **runner** is a headless Swarmfile engine that watches one branch and executes a job when a commit or tag matches a rule — the same shape as GitHub Actions or a self-hosted CI runner, but built on the engine's existing content-addressed blocks instead of a full `git clone`. Each rule match materializes the branch into a plain checkout directory (not a live mount) before running the job. By default that is every file in the tree — the checkout directory persists between runs, so a re-checkout only re-fetches blocks that actually changed since last time (an unchanged file is a fast local copy, not a network fetch), but the first run against a large project, or one where most files change often, still pays for the full tree. A top-level `sparse:` list narrows the checkout to the paths a job needs, so a huge repo only materializes the part it uses.

What it is not

This self-hosted runner has no job queue, no matrix builds, and no artifact storage — it is the trigger→execute→report loop only. For those, use Swarmfile's separate [Hosted CI](#), which runs jobs on Swarmfile-managed containers. Run history is inspectable via the CLI, the Desktop App, and the web dashboard (below).

A runner is also deliberately **not a mounted engine**: it serves no control socket and never presents a drive, so the `swarmfile` CLI cannot talk to it. It is configured entirely by environment variables plus the `.swarmfile/runner.yml` in the branch it watches, and a job starts only when a matching commit or tag lands — this runner has no external "run this now" API. On a machine that *does* run a mounted engine, `swarmfile tail` streams this project's activity in real time; the org's outbound webhooks are delivered on a periodic schedule, not a real-time build trigger; treat them as a reaction/integration channel.

The watched branch must be protected. `.swarmfile/runner.yml` decides what runs on the runner host, so a runner executes it only on a branch with effective [branch protection](#) — otherwise anyone who can push to that branch could run commands on the machine. `SWARMFILE_RUNNER_ALLOW_UNPROTECTED=1` overrides that and is not a production option — use it only on a trusted or throwaway runner.

For a job that isn't trigger-driven, a machine running a mounted engine can use `swarmfile materialize <ref> --out <dir>` — the same checkout, written to a plain directory with no FUSE, pinned to whichever ref you name.

Where rules live

Rules are **not** configured on the hub, in the Desktop App, or via the CLI — there is no `swarmfile runner rules` command, because there is nothing on the hub to create. A rule is a plain YAML file, committed into the branch it governs, at:

```
.swarmfile/runner.yml
```

This is deliberate: a rule change is reviewable in the same diff as the code it gates, exactly like a GitHub Actions workflow file. Whoever can commit to the branch controls what runs — the same trust boundary as any other file on that branch, not a new one.

```
# Optional, top-level: check out only these paths (see "Checkout scope").
sparse:
  - "src/**"
  - "*.md"
rules:
  - name: build-and-test
    on: commit          # commit | tag - default: commit
    branches:          # glob patterns; omit to match every branch
      - main
      - "release/*"
    paths:              # optional - matches the full relative path
      - "*.dwg"
    run: "scripts/ci.sh" # shell command, run with the checkout as CWD
    timeout: 1800       # seconds - default 1800 (30 min); on expiry the command's whole process tree is killed
```

A project can define several rules; each is evaluated independently against every commit/tag on the watched branch.

paths: matches each changed file's full relative path. `*.dwg` matches `plan.dwg` wherever it lives in the tree; `assets/**` anchors to that folder. A leading `!` negates (gitignore-style, last match wins), so `paths: ["src/**", "!src/docs/**"]` skips doc-only changes; the same applies to `branches:`. Leave `paths:` off to run on every commit to a matching branch.

on: tag rules ignore **paths:**. A tag names a single commit; there's no "files changed since the last tag" the way there is for a fresh commit.

Checkout scope (sparse)

By default a matched run checks out the whole branch. For a large repo, add a top-level `sparse:` list to materialize only the paths a job needs:

```
sparse:
  - "src/**"
  - "assets/models/*.bin"
  - "*.md"
rules:
  - name: build
    run: "make -C src"
```

A path is checked out when it matches a glob, or when a directory in its ancestry matches. A glob with no `/` matches by name at any depth (`*.md` finds `docs/notes.md`); a glob containing `/` is anchored to the repo root (`assets/models/*.bin` does **not** match `vendor/assets/models/x.bin`). A leading `!` excludes a path an earlier glob included (gitignore-style, last match wins), so `["src/**", "!src/**/*.test.ts"]` checks out `src/` without its tests; the top-level `exclude:` key is sugar for those `!` negations, and with no `sparse:` it means "everything except". Paths that match nothing are neither written nor fetched: a directory the filter cannot reach is skipped without being walked (and a frozen commit tree is not even fetched), and a directory holding no matching file is not created. Only **anchored** globs enable that pruning — a bare name like `*.md` can match at any depth, so the walk cannot skip anything; anchor with a `/` (`docs/*.md`) when you can.

`sparse:` is **top-level, not per-rule**: a runner has one checkout directory and materializes it once per event for every rule that matched, so a per-rule scope could not be honored. Omit it (or leave it empty) for a full checkout. Changing the list between runs converges the directory: paths that drop out of scope are removed, the newly included ones are fetched. At most 64 globs of 1024 characters each, and at least one non-`!` glob; an invalid list fails the affected runs with the reason (visible in `swarmfile runner runs`) rather than silently checking out everything.

Running a runner

A runner is a normal `swarmfile-engine` process started with `SWARMFILE_RUNNER_MODE=true` and a project API key (see [CLI: swarmfile](#) for `swarmfile api-key create` — the key counts against your plan's headless-key allowance; see [Billing & Plans](#)). Like seed mode, it never mounts a drive and never prompts for interactive sign-in; unlike seed mode, it materializes a real checkout directory (under its cache dir) because a job needs actual files on disk to run against — not a FUSE/WinFsp mount, which would require a kernel driver on a CI box that may not have one.

```
export SWARMFILE_API_KEY=sf_key_...
export SWARMFILE_PROJECT_ID=proj_xyz789
export SWARMFILE_BRANCH=main
export SWARMFILE_RUNNER_MODE=true
swarmfile-engine
```

`swarmfile-runner` is the same thing as its own binary: it sets `SWARMFILE_RUNNER_MODE=true` (when unset) and re-execs the sibling `swarmfile-engine`, so the CI role reads as its own tool without an env-var prefix. It ships with the installers — `/usr/bin/swarmfile-runner` via the Linux `.deb`, inside the macOS app bundle (with a `/usr/local/bin` symlink), and next to `swarmfile-engine.exe` from the Windows MSI. On any engine binary the env-var form above is equivalent. See [swarmfile-runner](#) for its flags and acquisition. It takes the same env and flags:

```
SWARMFILE_API_KEY=sf_key_... SWARMFILE_PROJECT_ID=proj_xyz789 swarmfile-runner
```

`SWARMFILE_BRANCH` is the same setting a normal branch-pinned mount already uses (see [Branches & Merging](#)) — "which branch does this runner watch" needs no setting of its own.

On startup the runner replays any commits it missed while it was offline (via the branch's commit log), then listens live. A crash or restart never silently skips a commit.

Before checking anything out, a triggered run waits up to 120 seconds for a file that is still uploading on the branch — a mid-change tree would be the wrong thing to build. If the upload is still live at the deadline, the run is not started: it is recorded as a `failure` whose log says it was skipped for a live upload (visible in `swarmfile runner runs`), rather than building from a partial tree. There is no separate `skipped` status. A shutdown during the wait ends it immediately, treated the same way.

Checking out a ref in your own CI

If you already have a CI system, you don't need runner mode to get a checkout: `swarmfile materialize <ref> --out <dir>` writes a ref's tree to a plain directory with **no mount and no FUSE**, so it works on a container with no kernel driver. Pin by **hash** — a tag is a name that can be repointed by delete + recreate, and a branch moves; the 64-hex hash is the only immutable pin (a commit's hash is shown in the web commit log, MR header, and release page, and on `swarmfile log / show`):

```
swarmfile materialize hash:$SWARMFILE_COMMIT_HASH --out "$PWD/checkout"
```

`--json` echoes the resolved `commitHash`, `kind`, and `seq` — record them in your build metadata so an artifact names the exact tree it was built from.

Add `--sparse <glob>` (repeatable) to check out only the paths a job needs — the same scoping as the runner's top-level `sparse:`, with the same glob rules:

```
swarmfile materialize hash:$SWARMFILE_COMMIT_HASH --out "$PWD/checkout" --sparse "src/**"
```

A few things to know before relying on it:

- **The executable bit is preserved.** A file marked executable (with `chmod +x` on the drive, `swarmfile chmod +x`, or a git push of mode `100755`) checks out executable on macOS and Linux, so `./scripts/ci.sh` runs as-is. Windows has no executable bit.
- **The marker owns the directory.** A new or empty destination is materialized fresh; one this tool already wrote is updated in place, pruning only the paths in its marker, so your job's own build outputs and caches survive a re-checkout. A non-empty destination with no marker, or one owned by a different project, is refused rather than merged or deleted — there is no `--prune` flag.
- **One destination, one writer.** The marker's written-paths list is read and rewritten per run with no cross-process lock. Give each job its own `--out` directory; two concurrent materializations into the same one can interleave.
- **Symlinks are written, not copied.** On Windows that needs Developer Mode (or elevation); a symlink the OS refuses is reported in the run's per-file `failures` rather than silently skipped, like any other entry that can't be written.
- **Not every commit has a hash.** A commit finalized before hashing was available, and an archived commit, have none; address those by `seq` or `tag`. An unknown hash is the typed `hash_not_found`, not a silent fallback.

`swarmfile clone <project-id> <path> --checkout -b <ref>` is the git-familiar spelling for the same operation, cross-project included: `-b` takes the full ref vocabulary (branch, tag, hash, seq), and `--sparse <glob>` (repeatable) narrows the checkout exactly as it does for `materialize`. The mount form of `clone` keeps `-b` branch-only, because a mount works one branch and a tag or hash has no branch to live on; it refuses `--sparse`, since a mount always shows the whole branch.

`materialize` still needs a running engine (the CLI is control-socket-only), so on a driverless box it's usually a mounted engine elsewhere doing the checkout, or the runner above doing the whole loop. The `/runner-runs` check-reporting API below needs no engine at all.

Inspecting run history

```
swarmfile runner runs          # every branch in the project
swarmfile runner runs --branch main  # narrowed to one branch
```

Each row carries the matched rule's name, the triggering commit, status (`running` / `success` / `failure` / `timed_out`), exit code, and a truncated log. There is no separate `swarmfile runner runs show <id>` command — the list output already includes each run's log inline, so there's nothing further to drill into.

The same history is browsable without the CLI: in the Desktop App, via the project left rail's **CI runs** row; on the web dashboard, under the project's **Runs** tab (open to every project member, like **Files** and **History** - guests are refused). Both surfaces also offer **Run workflow** and **Cancel** for a project's runs — those act on Swarmfile's hosted-CI service, never on a self-hosted runner, whose runs are only ever created and finished by the runner engine. The adjacent **CI** tab, which holds hosted-CI settings, secrets and runner sizes, is the admin/owner-only one (like **Commits**, **RfIs** and **Permissions**).

Reporting checks from your own CI system

You don't have to run Swarmfile's own runner to gate a `protected branch's --require-checks-to-pass` — any CI system that can make two authenticated HTTP calls can report a check result the same way a runner does. The `/runner-runs` API is deliberately not tied to the runner binary: it only requires a project API key, the same credential a runner engine uses (see `swarmfile api-key create`).

The two calls are the whole contract:

```
# Start of your job — creates a `running` row so a crash reports as failure,
# not silence.
RUN_ID=$(curl -sf -X POST "$SWARMFILE_HUB_URL/runner-runs" \
  -H "Authorization: Bearer $SWARMFILE_API_KEY" \
  -H "Content-Type: application/json" \
  -d '{
    "projectId": "'"$SWARMFILE_PROJECT_ID"'",
    "branchId": "'"$SWARMFILE_BRANCH_ID"'",
    "branchName": "'"$SWARMFILE_BRANCH"'",
    "changesetSeq": "'"$SWARMFILE_CHANGESET_SEQ"'",
    "ruleName": "external-lint"
  }' | jq -r '.run.id')

# ...your actual CI job runs here...

# End of your job — reports the real outcome.
curl -sf -X PATCH "$SWARMFILE_HUB_URL/runner-runs/$RUN_ID" \
  -H "Authorization: Bearer $SWARMFILE_API_KEY" \
  -H "Content-Type: application/json" \
  -d '{"status": "success"}' # or "failure" / "timed_out"
```

What each field means for gating, exactly as a Swarmfile-runner-produced row would be read: `ruleName` is an arbitrary label you choose — it's what shows up as one named check. `requireChecksToPass` first narrows to only the runs reported for the exact commit being merged, then requires the *latest* run of *every distinct* `ruleName` reported at that commit to be `success`. A rule that reported at an earlier commit but hasn't reported again at the new head isn't carried forward as still-required — it simply isn't part of the gate for this head, the same as a rule that's never run at all; a branch with zero matching runs at its head fails closed rather than passing vacuously. `changesetSeq` must be the commit number of the exact commit your job ran against (`swarmfile status /mr status` locally, or resolved from whatever your CI already checked out) — a check reported against an older commit doesn't satisfy a merge request whose branch has since moved past it. `branchId` is the branch's real id, not its name; resolve it once via `swarmfile branch list` or `GET /branches`.

A GitHub Actions job reporting into this looks like:

```

- name: Report check start
  run: |
    RUN_ID=$(curl -sf -X POST "$SWARMFILE_HUB_URL/runner-runs" \
      -H "Authorization: Bearer ${ secrets.SWARMFILE_API_KEY }}" \
      -H "Content-Type: application/json" \
      -d '{"projectId": "...", "branchId": "...", "branchName": "...", "changesetSeq": "...", "ruleName": "gh-actions-build"}' \
      | jq -r '.run.id')
    echo "RUN_ID=$RUN_ID" >> "$GITHUB_ENV"

# ...your build/test steps...

- name: Report check result
  if: always()
  run: |
    STATUS=${ job.status == 'success' && 'success' || 'failure' }
    curl -sf -X PATCH "$SWARMFILE_HUB_URL/runner-runs/$RUN_ID" \
      -H "Authorization: Bearer ${ secrets.SWARMFILE_API_KEY }}" \
      -H "Content-Type: application/json" \
      -d '{"status": "$STATUS"}'

```

This is exactly what Swarmfile's own runner does internally — there's no separate, more-privileged path only the runner binary can use. A project API key is the whole trust boundary, same as everywhere else in the API (see [Security](#) below).

Security

A rule's `run:` command runs with the permissions of that machine's user — the same trust model as a self-hosted GitHub Actions runner. The runner's API key itself is narrowly scoped: it can create and update its own run records, plus everything any other project-scoped API key can already do (read metadata, download blocks, read/write changesets and branches). It cannot post comments, touch ACLs, or reach admin routes — the same restriction every API key has today.

Hosted CI

Hosted CI runs the workflows in `.swarmfile/ci.yml` on Swarmfile-managed containers. There is nothing to install on a build machine and nothing to keep patched: a commit (or a manual dispatch) starts a container, the container streams the job's log back as it runs, and the result lands on the branch as a normal check.

Hosted CI is off for a project until an owner enables it, and a workflow only runs automatically from a branch with **effective protection** unless the project explicitly opts into unprotected branches (see [Security](#)). Every job is metered per second while it runs.

For a headless runner you host yourself, see [Runner as CI](#); the hosted service below shares the check and status vocabulary but not the machine.

Enable it

1. Open the project in the web app → **CI** tab (or the tray's **Project settings** → **Hosted CI**, or run `swarmfile ci enable`).
2. Toggle **Enable hosted CI for this project**. This is the cost guard: no workflow runs until an owner turns it on.
3. Commit a `.swarmfile/ci.yml` to the branch you want to run from.

The **CI** tab also holds the project's secrets and variables and shows the runner-size catalog (every built-in and the org's custom sizes, each with its per-second rate, read-only). Custom sizes are org-wide and are managed in **Settings** → **Hosted CI** (or `swarmfile ci size`). The **Runs** tab lists every hosted run on the branch, with a per-job log viewer and a **Cancel** button; like **Files** and **History** it is member-only — external collaborators (guests) are refused. A finished run can be re-run from its page (`swarmfile ci rerun <run-id>`), and a single failed job re-queued (`swarmfile ci retry <job-id>`). A run that ends in failure — including a timeout, a cancel, a queue expiry, a blocked start or an infrastructure fault — notifies its starter, and on a protected branch that branch's recent committers too, by bell/tray and the personal webhook, not email (see [Notifications](#)).

A first workflow

```
version: 1
name: Build and test

on:
  commit:
    branches: [main]
  manual: true

jobs:
  test:
    size: medium
    timeout: 20m
    steps:
      - name: Install
        run: npm ci
      - name: Test
        run: npm test
```

Commit that to a protected `main` and a run starts on the commit. The **Runs** tab shows it live; `swarmfile ci logs <job-id>` streams the same text in a terminal.

Triggers

Trigger	Runs when	Filters
<code>on.commit</code>	A commit lands	<code>branches</code> (globs, <code>!</code> negates), <code>paths</code> (changed paths; empty match means no run)
<code>on.tag</code>	A tag is created	<code>pattern</code> (globs)
<code>on.manual</code>	Someone dispatches it — the Runs tab's Run workflow , the tray, or <code>swarmfile ci run</code>	<code>inputs</code> (optional; see below)
<code>on.merge_request</code>	An MR opens, its head moves, it reopens or becomes ready for review	<code>types</code> , <code>branches</code> (base-branch globs), <code>paths</code>
<code>on.schedule</code>	A 5-field UTC cron expression fires	<code>crons</code> , <code>branches</code> (default branch if omitted)

A `paths` filter on a commit or an MR compares against the changed entries; if the change set is unknown the filter matches nothing (a run is never started on a guess). A schedule fires at most once per missed occurrence: a hub that was asleep catches up the **latest** due firing, not every one.

Manual inputs

A manual trigger can declare typed inputs. The dispatch form (web or tray) renders them automatically, and the CLI passes them with `swarmfile ci run --input name=value` (repeatable; values travel as strings and coerce to the declared type):

```

on:
  manual:
    inputs:
      environment:
        description: Where to deploy
        type: choice
        options: [staging, production]
        default: staging
      dry_run:
        type: boolean
        default: false
    retries:
      type: number
      required: true

```

Types are `string` (default), `boolean`, `number` and `choice` (with `options`). `required: true` with no default must be supplied; an optional input with no default gets the type's zero (`""`, `false`, `0`). Values are read as `${{ inputs.<name> }}` anywhere an expression is allowed — including a job `if:`, which is evaluated when the run is created, so a job can be skipped statically. An unknown input name, a missing required one, or a value that does not match the declared type (or choice options) refuses the dispatch with a visible `ci_input_invalid` blocked run. Input names use letters, digits and underscores.

Jobs and steps

```

jobs:
  build:
    name: Build
    size: large           # small | medium (default) | large | xlarge, or a custom size
    image: swarmfile-node # optional; default swarmfile-standard (see Runner images)
    environment: production # optional; binds the job to an org environment (see Environments)
    approval: true        # optional; force an approval even if the environment names none
    timeout: 30m           # default 30m; 120m maximum (Enterprise up to 360m)
    needs: [setup]         # waits for every listed job
    continue-on-error: true
    if: success()
    env:
      NODE_ENV: production
    secrets: [NPM_TOKEN]  # ONLY these secrets are readable by this job
    outputs: [artifact]
    sparse: ["src/**"]    # optional – check out only these paths
    steps:
      - id: prep
        name: Prepare
        run: echo "sha=$(git rev-parse HEAD)" >> "$CI_OUTPUT"
      - name: Publish
        if: success()
        working-directory: packages/app
        shell: bash
        env:
          TOKEN: ${ secrets.NPM_TOKEN }
        run: npm publish --provenance

```

- Every job runs in its own fresh container with a clean checkout of the triggering commit.
- `run:` is a shell script; the default shell is `bash -e -u -o pipefail`, overridable per step.
- `if:` on a job or a step accepts `success()`, `failure()`, `always()`, `cancelled()`, comparisons, `!`, `&&`, `||`, parentheses, and the functions `contains`, `startsWith`, `endsWith`, `format`, `join`, `toJSON`, `fromJSON`. A condition the runner cannot evaluate is treated as **false** — a job never runs on a condition no one can explain.
- `needs` jobs that did not succeed cause their dependents to be skipped (`ci_needs_failed`) — unless the dependent carries an explicit `if:` other than a literal `success()`, which is admitted and evaluated with `failure()` true for the failed need. A job with `continue-on-error: true` counts as success for `needs`.
- `sparse:` checks out only the listed paths, so a huge repo only materializes what the job uses — see [Partial checkout](#).

Partial checkout

Every job gets a clean checkout of the triggering commit. For a large repo, `sparse:` narrows it to the paths a job needs:

```

jobs:
  test:
    sparse:
      - "src/**"
      - "*.md"
    steps:
      - run: npm test

```

A path is checked out when it matches a glob, or when a directory in its ancestry matches. A glob with no `/` matches by name at any depth (`*.md` finds `docs/notes.md`); one containing `/` is anchored to the repo root (`assets/models/*.bin` does **not** match `vendor/assets/models/x.bin`). Everything else is neither fetched nor written, and a directory holding no matching file is not created. A leading `!` excludes a path that an earlier glob included (gitignore-style, last match wins; at least one non-`!` glob is required), and `#` comments are not supported. `exclude: [glob]` is sugar for those `!` negations — and with no `sparse:` it means "everything except" (check out the whole tree, then drop the exclusions). At most 64 globs of 1024 characters each. Omit `sparse: /exclude:` (or leave them empty) for a full checkout. The same globs work for the self-hosted `runner` and `swarmfile` `materialize --sparse`. Importing a GitHub workflow? An `actions/checkout` step's `sparse-checkout` becomes this field.

Matrix jobs

```
jobs:
  test:
    matrix:
      os: [linux, windows]
      node: ["18", "20"]
    steps:
      - run: echo "${{ matrix.os }}" on node "${{ matrix.node }}"
```

Each combination becomes its own job (`test (os=linux, node=18)`); a job that needs a matrix waits for all of its rows. `include:` adds or augments combinations and `exclude:` removes them, with GitHub's semantics. At most 64 combinations per job, 200 declared jobs per workflow and 256 expanded jobs per run — past the combination cap the workflow is refused `ci_matrix_too_large`; past a job cap, `ci_too_many_jobs` (declared) or `ci_jobs_too_many` (expanded).

Expressions and environment

`${{ ... }}` works in `env:`, `run:`, `working-directory` and `if:`. **Never interpolate an expression directly into `run:` for a value a teammate can influence** (`${{ inputs.* }}`, `ci.branch`, `ci.actor`, `vars.*`): like GitHub, the value is substituted into the shell command verbatim, so `run: ./deploy.sh "${{ inputs.environment }}"` lets a dispatcher inject shell syntax. Pass it through `env: instead — env: { TARGET: "${{ inputs.environment }}" }` then `run: ./deploy.sh "$TARGET"`. Only `steps.*` outputs are substituted whole at run time; an expression that *combines* a step output with other values inside `run:` or `env:` is not evaluated (the literal text reaches the shell) — compute it in a step and reference the output. The context is:

Context	Contents
<code>ci.*</code>	<code>branch</code> , <code>sha</code> , <code>tag</code> , <code>actor</code> , <code>event</code> , <code>run_id</code> , <code>run_number</code> , <code>status</code>
<code>env.*</code>	job <code>env:</code> values in a job <code>if:</code> ; a step's own <code>env:</code> is not visible to that step's <code>if:</code> (declare the value at the job level to branch on it)
<code>vars.*</code>	project variables, then org variables (project wins)
<code>secrets.*</code>	the job's declared secrets only
<code>needs.<job>.outputs.*</code>	a completed dependency's <code>\$CI_OUTPUT</code> values
<code>steps.<id>.outputs.*</code>	an earlier step's <code>\$CI_OUTPUT</code> values
<code>matrix.*</code>	the job's combination values
<code>inputs.*</code>	the manual run's resolved <code>workflow_dispatch</code> inputs

The runner exports `CI=true`, `CI_BRANCH`, `CI_SHA`, `CI_TAG`, `CI_ACTOR`, `CI_RUN_ID`, `CI_RUN_NUMBER`, `CI_JOB_KEY`, `CI_JOB_ID`, `CI_WORKSPACE`, `CI_PROJECT_ID`, `CI_ORG_ID`, `CI_EVENT` and `CI_OUTPUT` (the `name=value` file for job outputs). Print `::add-mask::<value>` and the runner masks that value (and the job token) out of the log from that point on.

Composite actions

Shared logic can live in a local composite action: a directory under `.swarmfile/actions/<name>/` containing an `action.yml` with `inputs:` and `runs.steps` (run-only steps; no nested `uses`, no `timeout`). A job step uses it with a repo-relative path and passes inputs with `with:`:

```
# .swarmfile/actions/setup/action.yml
name: Setup
inputs:
  node:
    default: "20"
runs:
  using: composite
  steps:
    - run: echo "using node "${{ inputs.node }}"
```

```
# .swarmfile/ci.yml (jobs.<id>.steps)
- uses: ../swarmfile/actions/setup
  with:
    node: "22"           # a scalar, or one whole "${{ ... }}" expression
- uses: ../swarmfile/actions/setup
  with:
    node: "${{ matrix.node }}" # the outer expression survives
```

The action's steps are expanded into the job at compile time: every `${{ inputs.<name> }}` becomes the caller's `with` value (or the declared default; a required input with neither refuses the workflow). The caller's `env:`, `if:` and `continue-on-error:` apply to the expanded steps (the action's own step `env` wins on a name

collision). A missing action file refuses the workflow `ci_action_not_found`; remote/marketplace `uses:` stays refused. Because the whole expansion happens before the run is created, an action change takes effect on the next run that reads the branch — and `swarmfile ci run` (manual dispatch) recompiles the branch head like any other run.

Runner images

`image:` selects the container a job runs in. Omit it and the job gets `swarmfile-standard` (Debian bookworm with git, git-lfs, Node.js, Python 3 and a C/C++ toolchain). The curated catalog is available on every plan:

Image	Includes
<code>swarmfile-standard</code> (default)	git, git-lfs, Node.js, Python 3, curl, build-essential
<code>swarmfile-node</code>	the standard image with Node.js 22 LTS

An organization can also register its own images (**Settings → Hosted CI → Runner images**): the reference must be digest-pinned and pushed to the managed registry first — pushing is all that's needed; jobs target the digest directly, with no deploy. Org names resolve before catalog names, so an org image may deliberately shadow a catalog name. Using an org image requires the Pro plan or above — a job on a lower plan is refused with `ci_byo_image_requires_plan`; an unknown image refuses with `ci_image_unknown`, and a reference the platform cannot pull fails the job with `ci_image_pull_failed`.

Environments and approvals

An organization defines **environments** (Settings → Hosted CI → Environments): a name, an optional URL, and a reviewer rule — zero or more reviewers as org members and/or owner/admin/member roles. A job binds to one with `environment: <name>`; if the environment names reviewers, the job **waits for an approval** once its `needs` are satisfied. Nothing runs and nothing is billed while it waits; a reviewer (any one of the named users, or a member whose role satisfies a named role — `role:admin` admits owners too) approves or denies it from the run page, and the bell notifies every eligible reviewer. Denial fails the job `environment_denied` and skips its dependents; approval queues it immediately.

A job may also declare `approval: true` beside `environment:` — that adds a gate when the environment names no reviewers (owners and admins then approve). A declaration can only add protection: it never broadens or removes the environment's reviewer set. Each environment can also turn on **"the run's starter cannot self-approve"**, which refuses the person who triggered the run even if they are a listed reviewer. A job naming an environment the org has not defined is blocked with `ci_environment_unknown`; retrying a failed gated job asks for approval again rather than starting unapproved. Reviewer changes apply to requests already waiting: removing a reviewer stops that person deciding a pending request, and clearing the reviewer list (with no `approval:` on the job) un-gates it — any member can then advance it from the run page. Deleting an environment leaves a pending request's reviewer set in place, so it can still be decided. Reviewer user ids are checked against org membership when the environment is saved. An environment also has its own **secrets and variables scope** (Settings → Hosted CI → Environments → **Secrets & variables** on the environment's row): a job that declares the environment resolves those on top of the project and org scopes, and the environment wins on a name collision — the natural place for deploy credentials that must not be readable by a job running on a pull request.

Secrets and variables

Project secrets and variables live in the project's **CI** tab; org-scoped ones are available to every project and live in the dashboard's **Settings → Hosted CI**, together with the org's environments (each with its own secret/variable scope), its custom runner sizes, its runner images and its CI policy — the org-level enable/disable control and the **Force unprotected runs off** override. All of these are owner/admin-managed.

Resolution precedence is environment (only for a job declaring it) → project → org: the most specific scope wins on a name collision.

- Secret **values are write-only**: they are never shown again, in any surface.
- A job must list the secret names it uses in `secrets:`. Referencing an undeclared secret fails the workflow before the job starts.
- Org secrets are only resolvable from **private** projects (an environment's secrets inherit this: a public project's workflows see project-scope values only).
- End-to-end encrypted projects do not run hosted CI: the runner holds no E2E keys by design.

Artifacts and dependency caches

A job can publish files for later jobs, and restore/save a dependency cache between runs:

```

jobs:
  build:
    artifacts:
      - name: dist
        path: packages/app/dist
        if-no-files-found: error # warn (default) | error | ignore
    cache:
      key: deps-${{ ci.branch }}-${{ ci.sha }}
      paths: [.npm]
      restore-keys: [deps-${{ ci.branch }}-]
    steps:
      - run: npm ci
      - run: npm run build

  deploy:
    needs: [build]
    downloads: [dist] # fetched after checkout, before steps
    steps:
      - run: ./packages/app/dist/deploy.sh

```

- **Artifacts** are bundled (tar.gz) at job end — on success or failure, or gated by the entry's `if:` — and `downloads:` fetches each name from the newest successful producer in the job's transitive `needs` closure (a missing producer fails the job). Limits: 1 GiB per artifact, 5 GiB per run, 30-day retention; a `path` that climbs out of the workspace is refused.
- **Caches** restore after checkout and before the first step: the exact `key`, then the newest entry matching any `restore-keys` prefix; the save happens only when the job succeeds. Key supports `${{ }}` (no secrets). Limit: 5 GiB per entry, 30-day retention.
- The GitHub importer maps `actions/upload-artifact`, `actions/download-artifact` and `actions/cache` (or the split restore/save pair) onto these fields when their arguments are literal.

Runner sizes, billing and limits

Size	vCPU	Memory	Mapped container
small	0.25	1 GiB	basic
medium (default)	1	6 GiB	standard-2
large	2	8 GiB	standard-3
xlarge	4	12 GiB	standard-4

Orgs can define custom sizes (quarter-vCPU steps, up to the platform ceilings); each custom size rounds **up** to the smallest offered container, and its per-second rate follows that container. Sizes are org-wide; manage them in **Settings** → **Hosted CI** or with `swarmfile ci size` (the project CI tab lists them read-only, with rates). Define one with `swarmfile ci size set build-heavy --vcpu 2 --memory 8 --disk 16` (or the size form in Settings → Hosted CI), then ask for it by name in a job: `size: build-heavy`.

Billing is per second of container time at the mapped container's rate (at-cost × 1.1), with a 60-second minimum per job. The run list shows each job's billed seconds and charge. Compute counts against the org's **spend cap** — over the cap, a new job is refused with a blocked `spend_cap_reached` run and reads keep working. A Community org needs a payment method on file before its first job (an owner adds it in **Billing**) and runs under the platform's Community compute ceiling; platform ceilings apply.

What a job costs

Rates are the provider's list price for each container, recovered at cost × 1.1, at current provider pricing:

Size	vCPU / memory / disk	Per minute	Per hour
small	0.25 / 1 GiB / 4 GB	\$0.00051	\$0.03
medium (default)	1 / 6 GiB / 12 GB	\$0.00237	\$0.14
large	2 / 8 GiB / 16 GB	\$0.00403	\$0.24
xlarge	4 / 12 GiB / 20 GB	\$0.00735	\$0.44

A job bills `max(60 s, elapsed)` at its resolved size's rate; the run list shows each job's billed seconds and exact charge. If provider pricing moves, this table and the in-app rates move with it.

How this compares to GitHub-hosted runners. Our billing is per second, with no included-minute allowance on any plan. GitHub's is per minute, with partial minutes rounded up, and its private-repo plans include a monthly minute allowance (public-repo jobs are free) — below that allowance, staying on GitHub can be cheaper. Past it, the comparable Linux job costs less here: our `large` (2 vCPU / 8 GiB) is \$0.00403/min, while GitHub's published 2-core rate works out about 50% higher per paid minute, and per-second billing widens the gap on jobs that don't land on whole minutes — a 5m30s job bills 330 seconds here instead of six minutes there. GitHub's larger, macOS and GPU runners are premium-priced and aren't offered on hosted CI; those jobs belong on the [self-hosted runner](#). Verify GitHub's current rates at their [pricing calculator](#), and see the dated side-by-side on the [GitHub comparison page](#).

Other limits: a job's `size` must resolve (an unknown name fails the job rather than guessing); a job's log is capped at 50 MiB and a run's log at 200 MiB (output past a cap is dropped and the log is marked truncated — the earliest output is what is kept) and logs are retained for 30 days; concurrent jobs are capped by plan (Community 1, Starter 2, Pro 8, Enterprise 16, weighted by size).

Security

A workflow can read the project's secrets, so hosted CI treats it like code that runs with credentials:

- Automatic runs come only from branches with **effective protection** (a protected branch rule) — a push to an unprotected branch is refused with a visible `blocked` run.
- A project owner can opt into unprotected branches (the **Allow unprotected runs** switch in the CI tab, `swarmfile ci enable --allow-unprotected`). Only do this when everyone who can push to those branches is trusted like a branch protector.
- A blocked run still writes its check row as a failure, so a required check can never be skipped by breaking CI.
- Job tokens are scoped to one job, one project and the triggering ref, expire with the job, and are stripped from every step's environment; the runner also blocks a step from reading the engine's `/proc/<pid>/environ` and masks the token from logs.

What is not supported

The dialect is small on purpose and fails closed: anything it does not understand refuses the workflow with a line number.

- No marketplace/remote `uses:` (never arbitrary code); a LOCAL composite action under `./.swarmfile/actions/<name>` IS supported (see Composite actions).
- No reusable workflows (`workflow_call`), `permissions`, or service containers. `environment:` + `approval:` ARE supported (see above). Artifacts and dependency caches are native (see below).
- No `container:` (GitHub's per-job container) — use `image:` instead, which selects one of the curated runner images or an org-registered one; `size:` selects resources independently.
- `workflow_dispatch` inputs ARE supported (typed `string` / `boolean` / `number` / `choice`; see Manual inputs). A manual run otherwise takes the branch only.
- No merge queues, no fork PRs, no `pull_request_target`.
- No self-hosted labels (`self-hosted`) — use the [runner as CI](#) instead.

Going further

- `swarmfile ci list|status|logs|cancel|run` and the `secret/variable/size` commands: see the [CLI reference](#).
- The desktop tray shows the same runs, logs and settings per project.

Automating Swarmfile

Swarmfile is built to run without a person at a keyboard: a render-farm node, a CI job, a coding agent, or a seed box can mount, sync, and report on a project on its own. This page is the map — which credential to reach for, the surfaces that are meant to be scripted, and what is deliberately not a public API.

The credentials

You need to	Use	Scope
Run an engine where nobody can complete a browser sign-in (render node, CI runner, agent fleet, seed box)	Project API key (<code>sf_key_...</code>)	Exactly one project
Act as yourself from a script — provisioning, webhook management, anything your own role can reach	Personal access token (<code>sf_pat_...</code>)	Whatever your account can do, live
Push and pull large objects with the stock <code>git-lfs</code> client	A git-LFS credential — a project API key minted from the project's git-LFS tab	That project's LFS traffic
Work at a keyboard	Normal sign-in	—

Project API keys are org-owned machine identities. A key is walled off to one project: it cannot read or write anything else in the org, and it is never treated as an owner or admin, no matter who minted it. Mint, list, revoke, and set a per-key request ceiling with `swarmfile api-key` (admin or owner), or from the dashboard's **API keys** page. The dashboard page also gives a key its own monthly **spend budget** — a dollar ceiling below the org's [spend cap](#); enforcement uses the lower of the two, and a key over its budget has only its own writes paused. See [CLI: api-key](#).

Personal access tokens are not a separate identity — a PAT resolves to you, and every access check runs against your own current role, so it can never do more than you can right now and it dies with your access, with nothing separate to clean up. Mint one under **Settings → Access Tokens** (self-serve for any member; up to 20 active per member, per org, optionally expiring). See [Personal access tokens](#).

Both are shown in full exactly once at creation, both are revocable with immediate effect, and both are audited. Two more rules worth knowing:

- On a machine with both an API key and a leftover `SWARMFILE_OIDC_REFRESH_TOKEN`, the API key wins — and the engine logs the choice rather than picking quietly.
- If your org enforces an [authentication policy](#), project API keys are exempt (they're org-provisioned machine identities), while a PAT snapshots the factors and email assertion of the session that minted it — a token minted before you enrolled stops satisfying the policy when enforcement begins.

Running an engine headlessly

Mint a key for the project it serves, then hand it to the engine as `SWARMFILE_API_KEY`:

```
# On a machine where you're an org admin/owner:
swarmfile api-key create render-farm-04 --project proj_9f2a

# On the headless machine:
export SWARMFILE_API_KEY=sf_key_...
export SWARMFILE_ORG_ID=...
export SWARMFILE_PROJECT_ID=proj_9f2a
swarmfile-engine
```

Key minting is a hard plan allowance, not a soft one: Free 1, Starter 2 per seat (floor 5), Pro 5 per seat (floor 10), plus key packs. A mint past the cap is refused with `keys_cap`; revoking frees the slot immediately, and a downgrade blocks new mints without revoking keys you already have. See [Billing & Plans](#).

Pin a key's own rate limit when a farm node should burst above the plan default:

```
swarmfile api-key rate-limit <key-id> 1000
```

For the full per-agent recipe — a labeled mount per agent and a shared dependency cache — see [Coding Agents](#).

Scripting the CLI

- `--json` (global) makes the CLI print the engine's raw response — the intended path for scripts — and a usage error comes back as `{"code":"usage_error",...}` instead of prose.
- Exit codes:** `0` success, `1` error, `2` a refusal a script should act on (permission, quota, a conflict that needs a human), `3` a long operation detached. See the [CLI reference](#).
- Long operations** (`merge`, `commit`, `checkout`, `materialize`, `project delete`, and more) wait and print their result by default. `--no-wait`, or Ctrl-C, returns the operation id (exit `3`) and the work carries on in the engine; collect the outcome later with `swarmfile op` — even after an engine restart.

CI

There are two shapes, and they compose:

- Swarmfile's own runner** — a headless engine that watches a branch and runs a job on every matching commit or tag, from an in-place checkout with no clone and no FUSE. See [Runner \(Headless CI\)](#).
- Your existing CI reporting a check** — any system that can make two authenticated HTTP calls can gate a protected branch's **require checks to pass** rule through `/runner-runs`, using the same project API key a runner uses. The runner guide has the exact two-call contract: [Reporting checks from your own CI system](#).

A headless engine authenticating with an API key also serves `git clone` and `git push` through the remote helper, once git access is enabled for the project — see [Cloning a Project with Git](#).

Webhooks

Webhooks deliver a signed HTTP `POST` to a URL you control when something happens in the org — no polling. An **org-wide webhook** is owner-managed, retried, logged, and auto-disabled after repeated failure; a **personal webhook** is a fire-and-forget notification channel for your own account. Deliveries are made on a periodic schedule, not in real time.

The full envelope, signature verification, event kinds, and limits live in [Webhooks](#). The automation-relevant part: an API key cannot manage webhooks or any other org-wide admin surface — that's deliberate. Manage them from a script with a personal access token instead, which acts as you.

Rate limits

Request budgets are per-user and per-org, counted in requests rather than bytes, and content streaming isn't counted. A PAT shares its owner's per-user window; an API key gets its own window at the plan's per-user default, adjustable per key. The engine batches and paces bursts against the budget the hub reports — burst behavior is covered in [Coding Agents](#).

What isn't a public API

There is no published REST API reference and no multi-language SDK. The supported programmatic surfaces are the ones on this page: the `swarmfile` CLI (with `--json`), headless engines authenticated by API keys, the `/runner-runs` check API, and webhooks. A personal access token can reach other routes your role can reach, but nothing beyond those four carries a stability contract — treat anything else as internal and subject to change.

Handling the credentials

- Copy a secret once and store it in a secret manager or your CI's secret store — never in the repo, an image layer, or shared shell history.
- Prefer an API key over a personal refresh token on shared machines; revoke a key the moment a node is retired.
- Keys are org-owned: deleting a project revokes its keys with it, and every mint, revoke, and rate-limit change lands in the org's [Activity feed](#).

Where to go next

- CLI: `api-key` — mint, list, revoke, and cap keys.
- [Personal access tokens](#) — the script credential that acts as you.
- [Runner \(Headless CI\)](#) — the runner, external check reporting, and ref materialization.
- [Coding Agents](#) — one mount per agent and a shared cache.
- [Webhooks](#) — event kinds and signature verification.
- [Billing & Plans](#) — key allowances and packs.

Sharing & Collaboration

Swarmfile has two collaboration surfaces: tools for your own org members working inside a project (comments, watch, presence), and tools for getting a specific file or folder in front of someone who isn't a member at all (share links, external collaborators). This page covers both, plus how they interact.

Comments and mentions

Every file has its own comment thread. Post from the CLI with `swarmfile comment <path> -m "message"`, or list a file's thread with `swarmfile comments <path>`. The web dashboard has a full thread panel with @-mention autocomplete for tagging teammates.

Threads live-update in the dashboard — anyone with the file open sees new comments and mentions arrive without refreshing.

A comment can be a reply, not just a new top-level note — click **Reply** under any comment to keep a back-and-forth in one thread instead of scattering it across separate top-level posts. Replies are one level deep (no reply-to-a-reply); resolving is a thread-level action on the original comment, not on each reply individually.

On an image preview in the web dashboard, click **Add pinned comment**, then click a spot on the image to leave a comment anchored to that exact location — useful for pointing at a specific area of a render or drawing instead of describing it in words. Pinned comments show as small markers on the image and, everywhere else the thread appears (including the Desktop App), as a 📌 badge next to the comment. Pinning is web-only in this version; PDFs and other file types don't support it yet.

Right-click any file in the web dashboard's Files view and choose **Open comments** to jump straight to its thread, without navigating there through the file list first.

Watch

Watch is a targeted, opt-in subscription to a specific file or folder, distinct from the general "someone else changed a file" bell notification you get for activity across a project. You choose exactly what you want notified about.

Watching a folder is automatically recursive — you don't pick recursive vs. non-recursive, the server derives it from whether you watched a file or a folder.

- `swarmfile watch <path>` — start watching a file or folder
- `swarmfile unwatch <path>` — stop
- `swarmfile watches` — list everything you're currently watching

Watch is also available outside the CLI: the Desktop App's file list has a per-row watch toggle, and the web dashboard exposes watch/unwatch from the right-click context menu on a file or folder. That same menu also has **Open comments** (above), **View history** ([Trash](#), [History & Rollback](#)), **Request unlock** ([Working with Files](#)), and **Copy Swarmfile path / Copy path in drive** for grabbing a file's location without opening it.

A watched-file change reaches you through the notification bell and a Desktop App toast; how notifications are delivered, and which kinds can also be emailed, is covered in [Notifications & Inbox](#).

Presence

Presence shows who's actively editing what, in real time. Run `swarmfile presence` (or the shorter alias `swarmfile who`) to see it from the command line, or check the dashboard for the same view. Presence is tied to the entry lock lifecycle, not a free-running heartbeat — see [Working with Files](#) for how locking and presence relate.

Share links

A share link publishes a single file or folder to anyone with the URL, without giving them a Swarmfile account or a seat.

When you create a link you can add:

- A password, stored hashed and verified server-side — the file's contents are never exposed by the password mechanism alone, see below
- An expiry date
- A cap on total access count

The access-count cap is enforced by the hub on every tier, but the unit differs. For managed-key shares it counts **downloads**: each download is checked *and incremented* against the cap. On an **end-to-end encrypted** project the cap counts **distinct verified visitors**: the first open by a verified email consumes one access, and re-opens (or block-token re-mints) by that same address don't. The cap is a hard limit on distinct verified email addresses for E2E shares. Expiry, revocation, email verification, and the per-recipient block list are all fully enforced regardless of tier, and revoking or blocking now takes effect on the very next block request rather than at session expiry. (The share dialog hides the access-count field for E2E projects; an API-created limit still enforces as described.)

Every share link requires email verification, regardless of whether you've also set a password. Before a recipient can see any content — download, preview, or the comment thread — they have to enter their email and click a one-time magic link sent to it. This isn't optional and it isn't skippable by knowing the password: the password (if set) and the email verification are both required, independently. Once verified, a recipient gets a 30-day session, so a returning collaborator isn't re-verified on every visit.

This gives the share owner a real access log: every verified recipient, by email address, with when they accessed the link. You can block any individual recipient from that log at any time, and the block takes effect on their very next request — an already-open browser tab doesn't get to keep working.

Shares support inline preview for images, PDFs, and text files, plus an optional public comment thread scoped to the shared item.

External collaborators

External collaborators are named guest reviewers — an M&E client, an AEC consultant — invited to a specific folder with `read` or `comment` permission, without consuming a paid seat. See [Permissions](#) for how folder-level access grants work generally; external collaborators are layered on top of that same system, scoped to a

folder rather than a whole project.

What a guest gets:

- A read-only mount of the folder they were invited to. Writes are refused by the mount itself, not a UI restriction — a `read` -permission guest cannot write to the mount even by bypassing the dashboard entirely.
- An offline-first comment and upload queue in the Desktop App, for guests granted `comment_upload` — the tier that adds uploads to comment access.
- Instant revocation and a full audit trail on the org side, with a durable revocation sweep you can watch under [Settings → Access revocations](#).
- If the same person is invited to multiple orgs as a guest, they get a picker in the Desktop App to switch between them.

External collaborators are plan-gated for **private projects**: Starter includes 1 per seat, Pro includes 5 per seat, and Enterprise is unbounded by contract. Starter and Pro both allow billed overage up to 3x the included allowance before new invitations are hard-blocked — see [Billing & Plans](#) for how overage is metered and charged.

On public projects, collaborators are free and consume nothing: a public project is readable by anyone, so a `read`, `comment` or `comment_upload` grant on it never occupies a plan slot and never accrues overage — which is why a Free-plan org, whose projects are all public, can have guests. A project can also refuse new guests for itself (inheriting or overriding the org policy; existing grants are not revoked) — see the project's settings or `swarmfile project show`.

On a public project, guests can also arrive without an invitation: **public access requests** are on by default (turn them off in the **Project settings** tab), so any registered user can ask for comment access — or **contributor** access, the `comment_upload` tier — and an owner or admin approves it from there. Approved grants land on the project's `Contributions` folder; everything above about the guest experience applies unchanged. See [Public Projects & Org Profiles](#).

Worksharing (no plugin)

Revit, SolidWorks, and many CAD/BIM tools already take a native operating-system lock on the central model while someone has it open — they open the file for writing and *deny write-sharing* to everyone else. On a plain network drive that lock only means something on one machine. Swarmfile **honors and enforces that native lock across every Windows machine** on the drive, so two people can't open the same central model for editing at once — with **no add-in and no separate lock server**.

Because it works at the filesystem level, not inside the application, there's nothing to install into Revit or SolidWorks. The app takes the lock it always takes; Swarmfile makes it real across every Windows machine on your team. (This enforcement is Windows-only — see [Platforms](#).)

How it works

When a native application opens a file on the mounted drive **for writing while denying write-sharing** — an *exclusive open* — Swarmfile registers a short-lived claim on that file for your machine. While that claim is live:

- Anyone on **another machine** who opens the same file the same way is refused. On Windows the open fails with a sharing violation, which the application surfaces the way it always does — "the file is in use," "could not open for writing," or the app's own central-model-locked message.
- The claim is renewed automatically for as long as the file stays open, so a model left open all day — or overnight — keeps enforcing. When the app closes the file, the claim is released and the next person can open it.

This is the same model a native app uses locally, extended across the internet. You don't change how you work: open the central model, and if someone else has it, your application tells you — just like it would on a shared drive that actually enforced the lock.

Who has it open

When your open is refused, Swarmfile also records a **Recent issues** entry in the Desktop App naming the user and machine that holds the file, so you're not left guessing at a bare "file in use." That is separate from the dashboard's [Presence](#) view, which is driven by Swarmfile's own lock lifecycle and can disagree with the passive open signal (see [Working with Files](#)).

Platforms

Native exclusive-open enforcement is a Windows capability, because the "open for write, deny write-sharing" semantics it builds on are a Windows filesystem concept — and Revit and SolidWorks are Windows applications. A cross-machine exclusive open is enforced between Windows machines.

macOS and Linux mounts don't take part. An exclusive open on a Mac or Linux mount registers no claim, and a Mac or Linux machine opening a file that a Windows machine holds exclusively is not refused by this mechanism. If a file is edited from mixed platforms, use an explicit edit lock instead.

For teams whose work spans macOS and Linux — the linked assets around a model, or a Rhino/QGIS/IFC ecosystem — the cooperative [entry locks](#) below work on every OS. The byte-range lock API is cross-platform too, but only on projects set to the Revit-worksharing project type — chosen at creation or set later with `swarmfile project set-type revit_worksharing` (see layer 3 below).

When the hub is briefly unreachable

Worksharing claims are **hub-enforced**: a claim is registered with the service, and every machine checks it before allowing an exclusive open. Deployments that need an open refused when the hub is unreachable can enable that per machine with `SWARMFILE_WORKSHARING_FAIL_CLOSED=1` on the engine (see [Environment Variables](#)); a claim that can't reach the hub then **denies** the open instead.

The three layers of locking

Worksharing is the top layer of Swarmfile's locking; each layer answers a different need. From coarsest to finest:

1. **Native exclusive-open enforcement (this page).** Whole-file, automatic, no plugin, Windows-only. The app's own "open the central model" is the lock. Best for Revit/SolidWorks-style worksharing where opening the file *is* checking it out.
2. **Entry locks.** Two related kinds. The **automatic entry lock** is the 60-second, heartbeat-renewed claim an open file takes while an app has it — nothing to manage. An **explicit edit lock** is one you take yourself with `swarmfile lock` (7 days by default, up to 30, auto-renewed every 6 hours while your engine is online — offline, the lease simply counts down) or a `git lfs lock` (a fixed 30-day lease that nothing renews); both kinds have a "request unlock" flow (`swarmfile unlock-request` from the CLI). A [Pack & Go reservation](#) is a separate mechanism that reserves a whole folder or project rather than one file. Best when you want to reserve a file you aren't holding open in an app. On a project used as a [git-LFS](#) server the entry lock is the **same** server-side lock as a `git lfs lock`, so the two mutually exclude across surfaces — and files that live only in git (never written to the drive) get an equivalent path-keyed lock so `git lfs lock / unlock / locks` still work for them.
3. **Byte-range locks.** Element-level locking *within* a single file — the way a BIM tool locks just the elements one person is editing without locking everyone out of the model. This is a lower-level API a native-app plugin calls directly; the `swarmfile-brlock` CLI is the reference client. It is enabled only on projects set to the Revit-worksharing project type — at creation, or later from **Project settings** or `swarmfile project set-type revit_worksharing`; on any other project the engine refuses byte-range requests.

Most teams never touch layers 2 and 3 for worksharing — opening the model is enough.

What this is not

Worksharing enforcement is a **write lock**, not a merge system. It stops two people editing one central model at once; it does not merge two divergent copies of a model back together (native BIM/CAD merge is the application's job). For versioning the model over time — branching a design option, rolling back a bad save — see [Version Control](#) and [Branches & Merging](#).

RFIs & Submittals

RFIs and submittals are structured, numbered, auditable review workflows layered directly over your project files. They're the coordination loop AEC teams run constantly — *ask a formal question and get a formal answer* (an RFI), or *submit a document and get it stamped* (a submittal) — but there's nothing construction-specific about the machinery. It's a review workflow with a paper trail, and it works for any team that needs "someone submitted this, someone decided on it, and here's the record of who and when."

The workflow is built by composing the pieces the rest of Swarmfile already provides: an RFI or submittal has a folder of attached files, so those attachments get [previews](#), comment threads, version history, locking, and ACLs for free. What the workflow adds on top is state, assignment, due dates, numbering, and a decision of record.

Who can use them

The **RFIs** tab in the web dashboard is visible to **owners and admins**. Access to an individual RFI's contents is governed by the ACL on the folder it lives under, the same as any other file — so a member granted access to that part of the project can see and act on it even though the top-level tab is admin-facing.

RFIs vs. submittals

They share one engine but differ in intent and in the decisions available:

- An **RFI** asks a question. Its decision is an **answer** — or, if you need it, **Approved** / **Rejected** / **Revise & Resubmit** / a plain comment.
- A **submittal** puts a document up for review. Its decisions are the review stamps: **Approved as Rev A**, **Approved as Rev B**, **Approved as Noted**, **Rejected as Noted**, and **Revise & Resubmit**.

You pick the kind when you create the record.

Creating one

Two entry points:

- **From a folder** — right-click a folder in the Files view and choose **Create RFI from this folder**. The new-RFI form opens with that folder as the attachment location already set.
- **From the RFIs tab** — create one directly and point it at a folder.

On the form you set the title and question (or submittal description), the kind (RFI or submittal), the reviewers it's assigned to, a due date, and a priority (**low**, **normal**, **high**, or **urgent**). You can attach files, which land in the record's own folder and behave like any other project files. Save it as a **draft** to keep working, or **submit** it to open it and notify the reviewers.

Lifecycle and ball-in-court

An RFI moves through a defined set of states:

draft → **open** → **in review** → **answered** → **closed**. A record can be **voided** at any point.

Alongside the status, every record tracks **ball-in-court** — whose turn it is, the originator or the reviewer. Submitting a draft puts the ball with the reviewer. **Revise & Resubmit** always leaves the ball with the reviewer, on both kinds — it's not a terminal decision, and a submittal specifically also bumps its revision (a → b → c, capping at c). What flips the ball back to the originator differs by kind: for an **RFI**, only the **answer** decision does — **Approved**, **Rejected**, and a plain **comment** all leave the ball with the reviewer. For a **submittal**, all four terminal review stamps do — **Approved as Rev A**, **Approved as Rev B**, **Approved as Noted**, and **Rejected as Noted** each hand the ball back to the originator and move the record to *answered*, since each one closes out the review; the originator's own *closed* action is the only step left after that. Ball-in-court is what the due-date reminders key off, so it's worth knowing which decision types move it.

Making a decision

A reviewer responds with a **decision** — an answer or a review stamp from the sets above — plus a message, and optionally a signature. The **decision timeline is the record of account**: it's a distinct, append-only history, separate from both the comment thread (informal discussion) and from file changes (which land as [changesets](#)). When you need to show *what was decided, by whom, and when*, that timeline is the answer, and it's what an audit looks at.

Requesting revisions sends the RFI back to *in review* with the ball returned to the reviewer; an answer moves it to *answered*; the originator accepting closes it.

Numbering

Each project has its own RFI counter, and numbers are **gap-tolerant**: a record gets its number when it's created (drafts included), and voiding it leaves that number as a gap rather than renumbering everything after it. No number is ever silently reused — the same convention tools like Procore follow, so the numbers you cite in correspondence stay stable.

Attachments, comments, and previews

Because an RFI's attachments are ordinary project files under the hood, everything else in the product applies to them: thumbnails and inline preview in the detail view, comment threads with @-mentions, per-file version history, and locking. Marked-up PDFs, drawings, and photos attached to an RFI preview inline the same way they would anywhere else.

Notifications

RFI activity generates notifications so nothing sits waiting on someone who doesn't know it's their turn:

- **RFI created** and **assigned** — to the reviewers it's routed to.
- **Response posted** and **revisions requested** — to whoever the ball moves to.
- **Due soon** and **overdue** — to the current ball-in-court, driven by a daily sweep.
- **Closed** and **voided** — to everyone involved.

There's also a **ball-in-court digest** for a periodic summary of what's waiting on you. These reach you through the in-app bell, Desktop App toasts, and email (with the same per-kind preferences and quiet hours as everything else — see [Notifications & Inbox](#)).

External consultants

The person who answers an RFI or reviews a submittal is often outside your org — a consulting engineer, an architect of record. Two paths, both fully audit-trailed:

- A **"respond" share link** — for first contact with no account. The consultant opens the link, verifies their email once (no Swarmfile account is created), sees the RFI read-only, and posts a decision directly. The decision is recorded against the share with their verified email, so the response has a real trail even though there's no account behind it.
- A **full external collaborator** — when the consultant is engaged properly, invite them as a guest (they get their own account). Their decisions carry their own identity, for a complete audit trail.

An external consultant you invite as a guest counts against your plan's external-collaborator allowance, the same as any other guest; a respond-link consultant has no account and doesn't. See [Sharing & Collaboration](#) and [Billing & Plans](#).

Public project intake

A public project opens a filing window (on by default, per project) so any registered Swarmfile user — a client, a consultant, a reviewer — can raise an RFI from the project's public page without joining the organization (turn it off under **Project settings** → **Public RFI intake**). The filing lands in an auto-created `RFIs/` folder at the project root with a priority the submitter picks; the team sees the submitter's verified name and email (recorded from their sign-in, never taken from their text), owners and admins get a dedicated notification and email, and the submitter gets an email when a member answers. Submitters track and reply to their filings under **My requests**, and members can filter the RFI list by public origin. Non-member attachments aren't accepted yet — a filing can describe the issue, and a member attaches files afterward.

Finding them

RFIs and submittals are full-text searchable by title and body from the dashboard omnibox, alongside files, projects, and people — see [Search](#). A match jumps straight to the RFI.

Desktop App and CLI access

From the Desktop App

Reviewing and deciding is coordination work, not work you do inside the mounted drive, so there's no RFI panel in the Desktop App. What it does surface is **RFI notifications as toasts**, and a deep link that opens the RFI in your browser, so someone working in the Desktop App still hears about an RFI waiting on them.

From the CLI

Full command coverage — create, list, decide, and manage RFIs and submittals without leaving a terminal, the same workflow a build script or a headless render-farm node can drive:

```
swarmfile rfi create ./Projects/tower-a/foundations --title "Slab thickness at grid C4" --body "Drawing S-201 shows 300mm, spec calls for 350mm
- which governs?" --priority high --reviewer user_88
swarmfile rfi list --outstanding
swarmfile rfi status rfi_9f2a
swarmfile rfi decide rfi_9f2a answered -m "350mm per spec - drawing will be revised"
```

Unlike `mr`, RFI commands (other than `create / list / search / by-file`) take an RFI reference that you can give either way: the friendly display number (`RFI-003`, case-insensitive, resolved through the mounted project) or the internal `id` from `rfi create / rfi list`. `rfi status <ref>`'s exit code doubles as a CI gate the same way `mr status`'s does: `0` if `answered / closed`, `2` for every other state. See [the CLI reference](#) for the full command list.

Where to go next

- [Sharing & Collaboration](#) — comments, share links, and external collaborators, which RFIs build on.
- [Notifications & Inbox](#) — how the RFI alerts above are delivered and configured.
- [Version Control](#) — changesets, which is where an RFI's *file changes* (as opposed to its decisions) are recorded.

Public Projects & Org Profiles

A project can be **public**: its releases are browsable and streamable by anyone with the link, it can appear on [Explore](#), and its organization gets a public profile page. (Downloading a release file or the release archive needs a free account; browsing, previewing and streaming byte ranges don't — and attached release assets are served anonymously by design.) Nothing about this changes how the project works for your team — the drive, locks, history, and permissions are identical.

Public projects are for distributing things — release assets, datasets, model weights, sample content. A team can also open a public project to **access requests**: any registered Swarmfile user (no organization membership needed) can ask for comment access, contributor access, or an invitation to join the organization, and an owner or admin approves each request. Approved guest and contributor requests on a public project are **free**: they consume no collaborator allowance and are never billed, on any plan (including Free). That's separate from [share links and external collaborators](#), which are for inviting specific people into private work.

What "public" means

Two things are set when a project is created and don't change afterward:

- **Visibility.** `Public` or `Private`. The choice is permanent — there's no "make this project public later" switch. If you need a public copy of private work, publish a release into a new public project (the **Copy to my org** action on a public release copies it into a project in your own org — always a public one, since it copies a public release; a **private** destination takes the Desktop App's CLI copy instead, see [Publishing releases → Getting a copy](#)).
- **Storage tier.** Public projects store content on the plaintext tier, because anything the world can download has nothing left to encrypt at rest. A project that needs managed or end-to-end encryption at rest can't be public.

Where a public project shows up depends on its releases. Its card appears on the org profile as soon as the project is public, but [Explore](#) and the project's own page stay empty until the first release is published — there is simply nothing to browse yet. Archiving a project removes it from the listings until you unarchive it.

The project page

Every public project has a page at `/public/<org-slug>/<project-slug>`. It shows the release's README (rendered, with its file listing), every published release with download counts and sizes, and the contributors of the latest release. From here a visitor can:

- **Browse a release** and download individual files (a free account is required for downloads), or the whole release as an archive where the release is set up for it.
- **Follow releases** — a verified free account gets an email whenever a new release is published. No account is needed to subscribe to the **Atom feed** instead.
- **See provenance.** A card or release marked **commit-pinned** means the release names the content-addressed commit it was published from. You can check the project's signed commit chain yourself with [swarmfile-verify-history](#); the badge itself is the pointer, not a verification result.

Requesting access

A public project accepts **access requests** by default (the team can turn them off in the **Project settings** tab), so the project page shows a **Request access** button to signed-in visitors who aren't members. Anything below is refused until your email is verified.

- **Comment access** — read the project and comment on files. No paid seat.
- **Contributor access** — comment, plus upload new files to the project's `Contributions` folder. No paid seat; no overwrite or delete rights, and uploading needs the desktop app or the web file browser, not a raw upload elsewhere.
- **Membership** — ask to join the organization. Approval sends the ordinary organization invitation (which uses a seat); you accept it from the email.

One request per project can be pending at once, and there's a small per-organization limit on simultaneous pending requests. You can **withdraw** a pending request from the page or from [My requests](#), where every access request you've filed and its outcome are listed. An owner or admin approves or denies from the project's **Project settings** tab; you get an email either way. If a request is declined, a new request needs a short message explaining what changed. An approved membership whose invitation expires unaccepted shows as **Invitation expired** in My requests, and the project page offers to ask again.

Public pages show the requester nothing beyond what they already see while signed out; your name and verified email are shared with the project's owners and admins only as part of the request.

The organization profile

An organization with public projects gets a profile at `/public/<org-slug>`: its logo, bio, website and support contact, a **README**, its public projects sorted by recently updated, most downloaded, or name (pinned projects first), and a contributors strip merged from its public projects. Member names are not shown — only people whose commits are already public on the projects themselves.

Owners and admins edit this under **Settings → Organization → Public profile**:

- **Bio** — one line, shown under the organization name.
- **Website and support contact** — shown as links on the profile.
- **README** — markdown, rendered below the header. Use full `https://` links; relative file links aren't resolved there. The editor has a **Preview** toggle and a **View public page** link.
- **Pinned projects** — up to 12 public projects to feature at the top, in the order you pick.

Every edit to the public profile is recorded in the organization's [audit log](#), with the fields changed but not their contents.

Following an organization

The profile's **Follow releases** button subscribes one verified email address to every public project in the org. The same email rules as the project follow apply:

- If you follow both the organization and one of its projects, you get **one** email per release, not two.
- Every email carries a one-click unsubscribe link, and you can also unfollow from the profile while signed in.

- Prefer a reader? The profile links an **Atom feed** that carries the org's releases across all its public projects: `/feeds/<org-slug>/releases.xml`.

What visitors can see, and what they can't

Visible on public pages	Never shown
Release files and their download counts	Anything from a private project
The README each release ships	Member names, roles, or email addresses
Contributor display names and commit counts	Account details or avatars
Org bio, website, README, pinned projects	Storage configuration or bucket details
Follower count (a number, never a roster)	The follower list itself

Publishers can also keep specific paths out of a release entirely — see [Keeping paths out of a release](#).

Archiving a project or deleting a release removes it from the public pages, the org profile, and Explore; public pages are cached at the edge for a minute or two, so a removal can lag by that much.

Publishing Releases

A **release** is a tag of a public project, published for anyone to browse and download: a fixed, cached snapshot of exactly what that tag points at. Releases are how you distribute a dataset, a set of sample files, model weights, an asset pack, or a game build from Swarmfile, without anyone needing to join your organization.

This page is the publisher's walkthrough. [Public Projects & Org Profiles](#) covers what visitors see and how your organization's public page works.

Before you start

- **The project must be public.** Visibility is chosen when a project is created and never changes, and public projects use the unencrypted storage tier (anything the world can download has nothing left to encrypt at rest). On the Free plan every project is public; on a paid plan, choose **Public** when you create the project. A private project can't be published — copy the content into a new public project instead.
- **You need the right access.** On a protected-mode project, publishing needs project-wide admin access (org owners and the project's creator always have it). On an open-mode project, any member who can write the whole project can publish. You can't publish a tag that contains anything you're denied read on, and you need a verified email. [Permissions](#) has the details, including how a publish fails safe when permissions can't be checked.

1. Tag the commit

A release always publishes a **tag**, and a tag always names one commit. Commit the state you want to ship, then tag it:

```
swarmfile commit -m "Dataset v1.0"
swarmfile tag create v1.0
```

Tags are immutable pointers for the native `swarmfile tag` command: there's no "move this tag." To point a name somewhere else, delete the tag and create it again — or, if the project has git access, move a lightweight tag through the git remote (`git push -f origin <tag>; git push origin :refs/tags/<tag>` deletes it). Once a tag is published as a release, treat it as permanent: people may already have downloaded it, a tag can't be moved, and the release keeps serving until you unpublish it. Tagging also protects that commit from history clean-up for as long as the tag exists.

2. Publish it

Publish from the dashboard's **Releases** tab (project → **Releases** → **Publish release...**). The wizard walks through choosing the tag, checking exactly what will become public, and confirming — the preflight shows the files and total size, anything `.swarmfile/publicignore.yml` keeps out, and any tracked paths `.gitignore` / `.swarmfileignore` hide from the web file browser that will still publish. If you fixed `.swarmfile/publicignore.yml` after cutting the tag, the preflight says the tag's policy differs from the branch's current one — publish a new tag to apply the fix.

You can also publish from the Desktop App's **Releases** panel, or from the CLI — `swarmfile release preflight v1.0` shows exactly what would publish (files, hidden paths, ignore-file mismatches, a policy that differs from the branch) without publishing anything, and exits non-zero when the tag can't be published:

```
swarmfile release preflight v1.0
swarmfile release create v1.0
swarmfile release list          # shows pending → ready (or failed)
```

A new release starts **pending** while Swarmfile builds its download archive, then becomes **ready** (or **failed**, with a reason). Re-running `release create` on a tag that's already published does nothing, unless the earlier attempt failed in a way that can be retried. If a publish failed permanently, tag a new commit and publish that.

To take a release down, use **Unpublish** on the Releases tab, or `swarmfile release delete <id>` (the id comes from `release list`). Unpublishing is a real delete, needs the same access as publishing, and removes the release from your public pages and Explore within a minute or two of edge caching.

What a release includes

- **The whole tagged tree**, browsable file by file on the project's public page at `/public/<org-slug>/<project-slug>`. Anyone can browse, preview, and stream byte ranges of a release's files without an account — a visitor can scrub a large video or dataset before downloading anything. (On a [bring-your-own-storage](#) org, anonymous streaming from your bucket is off until an owner enables **Allow anonymous streaming from my bucket**.)
- **Downloads behind a free account.** Downloading individual files, or the whole release as a single `.tar` **archive**, needs a signed-in (free) Swarmfile account. From the CLI: `swarmfile release fetch --org <org> --project <project> --tag v1.0 --out release.tar`, or add `--path` for one file. Files marked executable (`swarmfile chmod +x`) extract executable from the archive (mode `0755`); every other file extracts as `0644`. Serving those downloads is free: egress is unlimited and unmetered on every plan, so a widely used release costs storage, not bandwidth. (Anonymous reads carry an abuse backstop — a platform-wide hourly ceiling that can refuse further anonymous reads under extreme volume; see [Telemetry & Data Collection](#).)
- **The README** the release ships, rendered as the landing view of the project page.
- **Generated release notes** summarizing what changed since the previous release.
- **Download counts** per release (they lag by up to about an hour), and the contributors of the release.
- **Attached assets.** A release can carry named binaries that are deliberately outside the project tree — build artifacts, installers, model checkpoints — served at `/public/<org-slug>/<project-slug>/assets/<name>` and listed in the release page's Assets section. From the CLI:

```
swarmfile release attach v1.0 dist/widgets-1.0-macos-arm64.tar.gz
swarmfile release assets v1.0
swarmfile release detach v1.0 widgets-1.0-macos-arm64.tar.gz
```

The dashboard's Releases tab manages the same set from **Assets** on a release: attach a file up to 100 MiB, or remove one. Larger files attach with the CLI only, which uploads them directly to the project's storage (up to 5 GiB). Re-attaching the same name replaces the asset.

Keeping paths out of a release

Some of a project's files shouldn't go to the public page even when the rest should — internal notes, key material, licensed assets, an unpublished cut. Commit a `.swarmfile/publicignore.yml` at the project root and list the paths to hide (and, optionally, set how platform-generated derived data is exposed):

```
version: 1
paths:
  - "internal/*"
  - "**/*.secret"
  - "!internal/public/"      # re-include a folder a pattern above hid
```

`paths` uses the same gitignore-style globs as `.swarmfileignore`: `*`, `**`, a trailing `/` for a folder, and a leading `!` to re-include — last matching line wins. As in git, a folder must be re-included before anything inside it can be.

`derived` sets the visibility of data the platform generates around a release rather than files you committed. Its `ci_logs` key sets the exposure of CI-run logs attached to a release (`member-only` by default; `public` exposes them to anonymous visitors):

```
derived:
  ci_logs: member-only  # or public
```

`member-only` (the default when the file or the key is absent) keeps those CI logs visible to members only; `public` exposes them on the public page alongside the release.

The file is read from the version committed on the tag being published, and applied as the release is built:

- Hidden paths are left out of the public release entirely — no browse entry, no raw stream or download, nothing in the archive, file search or Explore. They aren't merely blocked; they never enter it.
- The policy file itself is never published, since its globs name exactly the paths it hides.
- If the file is invalid, the publish **fails with the reason** instead of publishing everything. Fix it and publish a new release. (A release that fails because storage couldn't be read can simply be retried.)
- Changing the file affects **future publishes only**. A release that's already published keeps the contents it was built with — publish a new tag to apply a change.

Members see everything, hidden or not; this is only about what non-members can reach. And hiding a path is not a retroactive takedown: it keeps the path out of new releases, but copies already downloaded stay where they are.

Provenance: commit pinning

A release names the commit hash its tag points at, and the project page shows it with a **commit-pinned** badge. Because commit hashes chain every file's content and all the history before it, the hash pins exactly what was published. The release also carries a manifest digest signed with the hub's Ed25519 history key, and the public page recomputes that digest in your browser and checks the signature, without downloading anything. The green **Verified by Swarmfile** badge appears only when the site build pins the hub's key; otherwise a valid signature reads *Signature valid - signed with a key published by this server, not independently pinned*, and a digest or signature that doesn't match is called out explicitly. The commit-pinned badge itself is the pointer, not a proof: anyone can check the project's signed commit chain with [swarmfile-verify-history](#), as [Verifiable History](#) explains, and `swarmfile materialize hash:<commit-hash>` reproduces the exact tree on any machine with access to the project. Owners can edit the generated release notes from the release page (a per-release summary of up to 4,000 characters), and the edit survives re-publishes.

Being found: Explore, following, and feeds

- **Explore.** Every public project with at least one release is listed on [Explore](#), Swarmfile's public directory, with search over names, descriptions, topics, README text and file paths. There's no separate opt-in: public plus a release means listed. Until the first release, a public project's card shows on your org profile, but there's nothing to browse yet. Archiving a project removes it from the listings until you unarchive it.
- **Follow releases.** A visitor with a verified free account can follow a project (or your whole organization) and gets one email per new release, with a one-click unsubscribe.
- **Atom feed.** Your organization's releases across all its public projects are published as a feed at `/feeds/<org-slug>/releases.xml`, no account needed.

Getting a copy: fork into your own org

On a release's public page, **Copy to my org** copies that release into a new project in an organization where the visitor can create projects, then offers to open it in the Desktop App. It's a one-time snapshot: the copy doesn't track the original, and later releases don't flow into it. Small releases copy straight away; a large one runs as a background job with progress and a cancel button. The destination organization's plan limits and storage quota apply, and copies are rate-limited per person. The copy the web flow creates is a **public** project (it holds a public release); pasting that release into one of your **private** projects runs on your own machine instead, through the Desktop App's CLI — `swarmfile clone <org>/<project> --into <your-org>/<project>` (add `--tag <tag>` to pick a release) — which encrypts the files there before they leave your computer.

This is also the way to get a public copy of work: publish into a public project, or copy a release into a new project of your own.

Current limitations

- Forks are one-time snapshots: there's no linked fork, stars, or public comments.
- Public commit history isn't available: visitors see releases, not the project's commit log.
- A public `git clone` URL isn't offered. Members of a project with git access turned on can `git clone` it; anonymous visitors can't.

Troubleshooting

Something not behaving? Start with the Desktop App's **Diagnostics** panel, or run the standalone report:

```
swarmfile-doctor --json --report doctor-report.json
```

It checks the engine, DNS, the hub, the sign-in token, cloud storage, the peer-to-peer network, mDNS, the firewall, and the mount itself, and its exit code is a verdict (0 all passed, 1 a failure, 2 warnings only). Most problems below show up there by name, and its output is the first thing support will ask for. See [swarmfile-doctor](#) for every flag.

The drive isn't there

- **Windows default drive letter changed.** If the usual letter was taken at startup, Swarmfile picks the next free one — that is per-machine and runtime-determined. Read the actual location from the Desktop App's Status/rail entry rather than expecting a fixed letter.
- **The boot mount's drive letter is held by something else.** The failure card offers **Use another drive letter** — it forgets any explicit letter you set and reconnects, so Swarmfile picks a free one; no `config.json` edit or manual relaunch needed. Open `Another Drive...` stays available too, for working at a different letter (or an empty folder) right now, without restarting. If the letter comes from `SWARMFILE_MOUNT_POINT` on that machine, the action refuses and says so — change the environment variable instead. If every letter is taken, choose an empty folder instead — `auto` has nothing left to pick.
- **A second mount failed.** Windows cannot mount over a folder that already has files in it: pick an empty folder, or let `swarmfile mounts open auto` choose the next free drive letter. A drive letter something else holds is refused with a message naming what holds it. From the CLI, `swarmfile mounts list` shows every open mount and `swarmfile mounts open <path> --label "..."` opens another.
- **The mount was healthy and stopped answering.** The drive is mounted to *wait* rather than fail mid-write, so a stalled mount looks like a hung Finder/Explorer window instead of an error. Open **Settings → Diagnostics** and use **Repair drive** (not available on Linux); it reconnects in seconds without touching your queued uploads. On Windows, an administrator can have Swarmfile do this automatically when a drive operation stays stuck, with the `SWARMFILE_FS_CALLBACK_RESCUE_SECS` setting (off by default). The full background is in [If the drive stops responding](#).
- **macOS: the drive keeps dropping** (the Desktop App's "keeps disconnecting" card). A leftover mount from a killed session is the usual cause; the engine clears those before each mount, so **Repair drive** normally reconnects it. If Diagnostics still reports a stuck mount, restart the computer — a reinstall won't clear it. See [If the drive stops responding](#).
- **The drive has never appeared and the app says the config file is broken.** A syntax error in `config.json` drops the whole file, including `org_id / project_id`, so status names the file and the parse position instead of asking you to pick a project. Fix the JSON and restart Swarmfile — see [Engine Config File](#).
- **Access is still being granted.** A project you have just been added to (or your account's first device) shows its drive only once its encryption key reaches this machine — for an end-to-end-encrypted project that means another member's client, or your org's granter (about a five-minute cadence at worst), has wrapped the key to it — and if no grant was ever queued, your engine asks the hub to open one itself. A running engine watches for the grant on its own and mounts the drive as soon as it lands, so leave the app open and let it finish rather than reinstalling or re-adding the project; **Recent issues** in Diagnostics shows the wait. See [End-to-End Encryption](#).
- **macOS/Linux:** the mount is a folder, and on a machine that has never mounted anything the engine may still be finishing startup. Give it a moment, then check Diagnostics.

A file won't open, or opens very slowly

- **First open pays for the network.** Opening a file you've never read streams from a LAN peer if one has it, otherwise from the cloud; the Desktop App's Status shows peer counts, and `swarmfile status` prints the same. If peers exist but transfers never come from the LAN, run the doctor — the `mdns listen` and `host firewall` checks name the cause.
- **macOS is silently denying local network access.** On macOS 15+, the engine needs *Local Network* consent (System Settings → Privacy & Security → Local Network) to find peers on your LAN. The doctor's `macos local network` check says so explicitly when a LAN connection fails for that reason, and that row in **Settings → Diagnostics** offers **Open Local Network settings** to take you straight there.
- **Corporate Wi-Fi blocks mDNS.** Discovery is multicast; on a VLAN or Docker bridge it may be blocked even though the hub is reachable. Set an office grouping (`swarmfile office set <name> --lan-from-office`) so same-office peers are treated as LAN anyway — see [Network Requirements](#).
- **If it's a wait, not an error.** A version that has never finished uploading can be opened; reads that outrun the upload return a retryable "try again" rather than an I/O error, deliberately, so applications don't mistake a slow upload for a damaged file. For the part you actually need, ask for it first: `swarmfile fetch <path> --tail 209715200`, or follow a moving viewer with `--follow` — see [Working with Files](#).
- **Every file fails, not just one.** If opens across a whole folder or drive return I/O errors while the listing still looks fine, the engine treats it as a stale key or access problem: it notices the pattern and refreshes the project's key and scope in place, and the files start opening again within a minute or two — no restart needed. **Recent issues** (Status/Diagnostics) shows the automatic refresh while it happens. If it is still failing after a couple of minutes, restart Swarmfile; that is the last resort, not the first step.

Writes are refused, or an app reports a plain I/O error

A filesystem can't say "someone else changed this too", so blocked writes look like ordinary I/O errors to applications. **The Desktop App remembers why:** click the Status pill and read **"Why did my save fail?"** — it lists recent refusals with the file's path, the reason, and how many times it happened (the engine keeps the last 64). Match the cause:

- **Unresolved conflict** — the same file changed in two places. The file shows as read-only (macOS/Linux) and saves fail with "permission denied" ("access denied" on Windows) rather than an I/O error. `swarmfile conflicts` lists them; resolve with `swarmfile conflicts resolve <path> --auto` (or `--winner local|remote|keep-both`). See [Branches & Merging](#).
- **Someone else holds the file** — the refusal names who holds it and their machine; `swarmfile locks` shows who, and `swarmfile unlock-request create <path>` asks them to release it. See [Entry locks](#).
- **A Pack & Go lease lapsed while you were offline** — the reservation expired (leases only renew while online). Re-run `swarmfile offline prepare` when you have connectivity; see [Working Offline \(Pack & Go\)](#).

- **A bulk Pack & Go reservation covers the file** — a single-file lock and a `scope` reservation both refuse writes with `locked`, but a scope reservation blocks checkout/restore/merge too, not just saves. Its owner releases it with `swarmfile offline return`; see [Working Offline \(Pack & Go\)](#).
- **A native worksharing tool has the file open elsewhere** — the Desktop App notification names who and which machine. See [Worksharing](#).
- **The file is still being created** (`pending_create`) — a lock, comment, share, or ACL action on a file whose first upload hasn't reached the hub yet. If the file is your own, renaming or deleting it from the app still works while it's pending (applied on this computer); this refusal is for the actions that need the hub. It's a hub `409` rather than a filesystem refusal, so it won't appear in "Why did my save fail?"; the Desktop App badge says "still being created". Wait a moment and retry, or check `swarmfile status → pendingCreates`.
- **The save is a new file at the project's top level** (`root_create_restricted`) — a protected project only lets its owner, its creator, or someone with project-wide write access add files at the root. The save is held and retried, not lost: ask an org owner to give you project-wide write, or create the file inside a folder instead. The same rule refuses a `git push` that adds a name at the root — the push error names the remedy; push the commit inside a folder instead. See [Permissions](#) and [Clone a project → Pushing](#).
- **You're offline and the file isn't in your reserved scope** (`offline`) — the drive is offline and this path wasn't covered by an offline reservation. Reconnect, or include it the next time you run `swarmfile offline prepare`.
- **This computer is quarantined** (`quarantined`) — the account was suspended after a burst of changes that looked like ransomware; an org owner clears it from **Settings → Quarantine**. On macOS and Linux, reads keep working and only writes are refused; on Windows the drive stays mounted and saves are accepted locally but simply don't sync until it's cleared. Already-queued saves are held, not lost — they resume once it's cleared. See [Security](#).
- **Another machine is still uploading the file** (`uploading_elsewhere`) — its first upload is still in flight from wherever it was created; wait for that to finish.
- **Windows lost a cached write** (`unflushed_write_lost`) — a save was refused. Reopen the file and save again (or Save As).
- **Another drive on this computer has the file open** (`locked_by_mount`) — with [exclusive-write mode](#), a drive refuses a sibling drive's write — and its rename, delete, or replace of that file — while it holds it (on Windows the write-open itself fails as "file in use"); the reason names the holding drive. Close the file there, or wait for that drive's save to finish.
- **Folder not empty** — a plain (non-recursive) delete of a folder that still contains files is refused, as on any disk. Delete it from your file manager, or with `rmdir /s` or `rm -r`, which empty it first. (PowerShell's `Remove-Item -Recurse` doesn't work through the mount — see [Filesystem Compatibility](#).)
- **Pack & Go disabled by policy** — an administrator can disable bulk offline reservation with `SWARFFILE_PACK_AND_GO_POLICY=disabled` on the engine; the CLI reports `pack_and_go_disabled`. If the value was meant to be different, check its spelling: an unrecognized value disables Pack & Go and now warns at startup and in `swarmfile-doctor`.

Files aren't syncing

The Desktop App header always reports the truth: **Syncing N**, **Synced**, or one of the attention states below.

- **Waiting on plan** — the hub is refusing saves on plan grounds, usually because the organization's storage allowance is full. Raise the limit or free space; retries resume automatically. The plan panel shows the allowance, and if the pause is the org's [spend cap](#) rather than the plan allowance, an owner or admin raises it under **Billing → Spend cap**.
- **N changes not syncing** — the hub keeps refusing these saves for another reason. They're safe on this computer and Swarmfile keeps trying. `swarmfile sync stuck` lists each file, why, and since when; `swarmfile sync discard <entry-id>` drops the local change and returns the file to its cloud version (it asks first, and it cannot be undone).
- **This computer is quarantined** — the account was blocked after a burst of changes that looked like ransomware. On Windows the drive stays mounted and saves are accepted locally but don't sync; on macOS and Linux writes are refused. An org owner or admin clears it from **Settings → Quarantine**, and queued work resumes on its own. See [Ransomware quarantine](#).
- **N changes failed** — work stopped retrying. The banner offers **Retry**; saving the file again also works.
- **A file was never supposed to sync** — check the exclusion rules: `swarmfile check-ignore <path>` says whether a path is ignored and by which `.gitignore/.swarmfileignore` rule. To remove already-synced, now-ignored content, `swarmfile ignore-clean` (soft delete, recoverable from Trash).
- **Signed out** — a lapsed sign-in pauses the queue without burning retries; sign back in and it drains. The doctor's `oidc token check` reports it directly.

A file others can see is missing for you

If a file others can see is missing on your drive, run `swarmfile cache resync`. It re-fetches this project's change feed from the beginning; nothing local is dropped and cached bytes are kept. Run `swarmfile doctor` to confirm the hub and the mount are healthy, and if the file still doesn't appear, [get help](#) with the doctor output.

Uploads are slow (very large file)

Saving never waits for the upload: the write returns at local-disk speed and the durable queue drains in the background. A multi-hundred-gigabyte save on a modest uplink can legitimately take days. To see what's live, `swarmfile uploads` lists every in-flight upload on the mount (yours and other machines'); `--wait` blocks until none remain, which is handy before a checkout in CI. While a file is uploading, other people can fetch the part they need first — see [While a large file is still uploading](#).

The same is true of **many files at once**: a mass edit or overwrite of hundreds of small files is paced to your plan's request allowance and can take minutes to reach **Synced** — see [Uploads are designed not to block](#).

Storage says full

Storage is enforced per organization, and a Free plan's allowance is a hard stop rather than a bill. Free space (`swarmfile hydrate free-space`, or **Free up space** in the Desktop App) or move the organization to a paid plan; if the pause is the org's [spend cap](#), raise it under **Billing → Spend cap** instead. The [Billing & Plans](#) page explains the allowance, overage, and where to change it.

Search returns nothing

Search queries a local metadata cache first, then the hub, so a file you've never seen in this project may not appear until the hub is reachable — and it matches names, not full file contents. See [Search](#).

An update didn't happen

Updates are never installed silently. The Desktop App checks for a new version at startup and periodically, and a found update only shows as a pill in the header plus a card in Diagnostics; **Install now** is the only path that installs it. If the pill has sat uninstalled for a week, you'll also see a one-time desktop notification reminding you — still no automatic install. If an update won't take, use **Diagnostics → Reinstall...** for a full reinstall that keeps your queued uploads. On Windows, the doctor's `updater-helper-task` check reports whether the updater helper is registered; `swarmfile-doctor --repair` does the same full reinstall headlessly.

Windows showed a SmartScreen warning during install

The blue **"Windows protected your PC"** box is expected the first time you run a newly downloaded installer — the `Setup.exe / .msi` installers are Authenticode-signed, but SmartScreen warns until their signing identity has built up download reputation, and that is not a sign anything is wrong with the download. Click **More info** (the small link, not **Don't run**), then **Run anyway** and finish the install; [Getting Started](#) has the step-by-step. It only appears when you run a freshly downloaded installer — the app's automatic updates don't show it.

If your IT department enforces an AppLocker or Software Restriction Policy, it may block Swarmfile's Explorer overlay DLL even though the app installs and syncs: files work, but status icons do not appear in Explorer. Ask IT to allow `C:\Program Files\Swarmfile`.

Signing in

- A message like "You may need to sign in again" means the session's refresh token expired or was revoked — sign in again and queued work resumes.
- SSO: enter your **organization's slug** (the short name in your org's Swarmfile URL). There's no email-domain auto-detection, so a personal email alone won't route to your IdP.
- A machine that can't sign in at all can still be diagnosed: `swarmfile-doctor` runs without an engine or a session.

"A file is still uploading" when branching or merging

`branch create`, `merge`, checking out or locking a file, and copying one are refused with `entry_uploading` while a file they'd pick up is still uploading. Taken mid-upload, they would silently use the file's previous version and still look complete. Usually you just wait for the upload to finish and retry. `swarmfile status` on the machine doing the writing shows its queue.

If that machine crashed or went offline mid-upload, restart Swarmfile on it: its queue picks up where it stopped and finishes the upload. If the machine can't come back, the stuck upload stops blocking branch creates, merges from its branch, and copies once it has sent no progress report for 30 minutes. Pausing sync doesn't count as silence: a paused machine keeps reporting, so its uploads still hold. Locking or checking out that file, and a merge that would change the file itself, keep waiting for up to six hours: taking them early would refuse the upload if its machine comes back. To clear the pointer explicitly — so readers stop waiting on it too — `swarmfile uploads retract <path>` (needs write access to the file; retracting a live upload is a no-op).

A merge won't finish

A large merge applies in the background once the hub accepts it — the command or app reports success before every file has landed, and `swarmfile merge --wait` blocks until it finishes. If one looks stuck:

- **Check the doctor.** The `merges in progress` check lists in-progress merges with their age; the service rolls back abandoned merges.
- **stale_merge is not stuck.** That refusal means the target branch moved after the merge was prepared; run `merge` again. It either lands cleanly or comes back with the genuine conflict that race exposed, which the normal resolution flow handles.
- **A project owner can repair a genuinely stuck one.** A repair action for owners and support resumes a merge that stopped, or rolls it back (forced rollback is owner/admin-only). Nothing is lost either way: entries already applied stay, and a follow-up merge picks up whatever is still divergent.

Getting help

When you contact support, include:

- `swarmfile-doctor --json --report doctor-report.json` (attach the JSON file).
- What you did, what you expected, and what happened — with the exact time and timezone.
- Whether it happens on one machine or everyone's, and on one file/folder or everywhere.

See [Operations](#) for admin-side diagnostics and connectivity checks, and [Network Requirements](#) if your firewall or proxy team needs the full list of hosts and ports.

Uninstall & Clean Reinstall

Removing Swarmfile from a machine is straightforward, with one important distinction: **repairing** an install and **removing** one are different operations, and repair is almost always what you want. If the app is misbehaving — a broken mount, a missing CLI tool, a half-applied update — use **Diagnostics → Reinstall...** or `swarmfile-doctor --repair --yes` instead. A repair reinstalls the full package and **keeps** your cache, queued uploads, config, and sign-in; removing the app and its data does not. A repair stops the engine briefly while the installer runs (on Windows the drive is detached cleanly first) and it reconnects when the install finishes; queued work is unaffected.

Everything below is for a deliberate removal — decommissioning a laptop, handing it to someone else, or a truly clean slate.

Before you remove anything

Three things live only on this machine and are lost when you delete its data:

1. **Queued uploads.** Saves that haven't reached the hub yet live in the local database. Check `swarmfile status` (the pending uploads / commits pending counters) and `swarmfile uploads --wait`; wait for zero, or accept losing those changes.
2. **Stuck saves.** `swarmfile sync stuck` lists saves the hub keeps refusing. Export or resolve them before wiping the cache — they will not come back.
3. **Reservations and locks.** `swarmfile offline return` releases any Pack & Go reservations this machine holds; checkout locks you took expire on their own lease, but releasing them now is neighborly. If this machine runs a **seed node**, disable seed mode first (`swarmfile seed disable`) so the org's discovery list doesn't keep advertising it.

Anything staged in a changelist that has been **parked** is kept on the hub and can be resumed from another machine; unsent changelist work is local and goes with the cache.

One more piece of wiring is separate from the app install: if this machine used a project as a [git-LFS](#) server, run `swarmfile-lfs uninstall` to remove the custom-transfer adapter from your git configuration. Uninstalling the app does not do that for you.

Windows

1. **Quit the Desktop App** from the tray (including any "quit engine" option so the mount dismounts). The uninstaller stops the engine and tray itself (with a force-kill backstop), but quitting first avoids holding the mount open. If uploads are still in flight, Swarmfile drains them before stopping; after about 30 seconds it asks whether to abandon the remaining saves or keep waiting.
2. **Uninstall from Settings → Apps → Swarmfile → Uninstall.** For scripted removal, find the product code and call `msiexec` (the CI harness does the same):

```
$key = Get-ChildItem 'HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall',  
                  'HKLM:\SOFTWARE\Wow6432Node\Microsoft\Windows\CurrentVersion\Uninstall' |  
    Where-Object { $_.GetValue('DisplayName') -match 'Swarmfile' } | Select-Object -First 1  
msiexec /x $key.PSChildName /qn /norestart
```

If the install came from the bundled `Setup.exe`, its uninstaller lives in `ProgramData\Package Cache`; Settings → Apps invokes it with `/uninstall` for you.

3. The uninstaller removes the program files (`C:\Program Files\Swarmfile`), the `PATH` entry, the `swarmfile://` protocol registration, the inbound UDP firewall rules, the Explorer overlay handlers, the Start Menu shortcut, the per-machine `Run` entries, and the **Swarmfile Updater** scheduled task. WinFsp, the VC++ Redistributable, and WebView2 are chain prerequisites and deliberately stay installed.
4. **Per-user data is deliberately left behind** (uninstalling for one user shouldn't wipe another's cache):

```
Remove-Item -Recurse -Force "$env:LOCALAPPDATA\swarmfile"
```

This removes the cache, queued uploads, config, and sign-in for that user only. One leftover is harmless: `HKCU\Software\Swarmfile\MountPoint` (the remembered drive letter), which the MSI doesn't own.

To reinstall, run the bundled `Setup.exe` again — or a **repair** if you only wanted the app fixed. If you downloaded a fresh `Setup.exe`, expect the SmartScreen prompt again; [Getting Started](#) walks through it.

macOS

The `.pkg` install is machine-wide: the app in `/Applications`, LaunchAgents in `/Library/LaunchAgents`, and CLI symlinks in `/usr/local/bin`. Removing it takes a few commands; run the `launchctl` lines as the logged-in user and the `rm` lines with `sudo`.

1. Quit the Desktop App, then stop its background services:

```
launchctl bootout gui/$UID/dev.swarmfile.tray 2>/dev/null  
launchctl bootout gui/$UID/dev.swarmfile.engine 2>/dev/null
```

2. Remove the app, the agents, the CLI links, and the bundled FUSE-T pieces:

```

sudo rm -rf /Applications/Swarmfile.app
sudo rm -f /Library/LaunchAgents/dev.swarmfile.engine.plist \
/Library/LaunchAgents/dev.swarmfile.tray.plist
sudo rm -f /usr/local/bin/swarmfile /usr/local/bin/swarmfile-*
sudo rm -f /usr/local/lib/libfuse3.4.dylib /usr/local/lib/libfuse3.dylib \
/usr/local/lib/libfuse-t-1.2.7.dylib /usr/local/lib/libfuse-t.dylib
sudo rm -rf "/Library/Application Support/fuse-t"
sudo pkgutil --forget dev.swarmfile

```

3. Remove the per-user data (cache, queue, config):

```
rm -rf ~/Library/Caches/swarmfile ~/.cache/swarmfile ~/.config/swarmfile
```

The engine's cache and config are the `~/Library/Caches/swarmfile` pair; `~/Library/Caches/swarmfile` covers `.dmg` and other install layouts. Repeat per user profile.

4. Remove the engine's macOS Keychain items — the sign-in token and this device's end-to-end key live there, not in the files above. Each item's service is `com.swarmfile.engine`; open **Keychain Access**, search for that string, and delete every match (or `security delete-generic-password -s com.swarmfile.engine`, repeated for each item). Do this per user profile too; the login keychain is per user.

If Finder or an app still shows the drive, it is a stale mount from a killed helper — log out and back in, or reboot, before wiping. `swarmfile-doctor --repair --yes` is the supported way to fix an install you intend to keep.

Linux

1. **Quit the Desktop App and the engine for each logged-in user first** — the package's pre-remove script disables the systemd units but does not stop services already running in your session, so a removed engine can keep serving its mount until you log out or kill it.

2. **Remove the package:**

```

sudo apt remove swarmfile      # remove the app and units
sudo apt purge swarmfile       # also remove the packaged /etc/swarmfile

```

Removal disables the per-user `swarmfile-engine` and `swarmfile-tray` systemd units; nothing else (firewall, mounts) is touched. If a mount is still present after removal, unmount it (`fusermount -u <mount point>`) or log out.

3. **Remove the per-user data** for each user who signed in:

```
rm -rf ~/.cache/swarmfile ~/.config/swarmfile
```

That covers the cache, queue, and config. If you ran a seed or a headless engine with its own `SWARMFILE_CACHE_DIR`, remove that directory too.

4. **Remove the keyring entries.** Where a desktop keyring is available (GNOME Keyring or KWallet), the sign-in token and this device's end-to-end key live there under the service name `com.swarmfile.engine` — delete those items in **Passwords and Keys** (Seahorse) or KWallet. On a headless/seed machine with no keyring, the engine falls back to a file in the cache directory, which step 3 already removed.

What lives where

Platform	Application	Per-user data (cache, queue, config, sign-in)
Windows	<code>C:\Program Files\Swarmfile</code>	<code>%LOCALAPPDATA%\swarmfile</code> (DPAPI secret store included)
macOS	<code>/Applications/Swarmfile.app</code> , <code>/Library/LaunchAgents</code> , <code>/usr/local/bin</code> , <code>/usr/local/lib/libfuse*</code> , <code>/Library/Application Support/fuse-t</code>	<code>~/Library/Caches/swarmfile</code> (cache, queue, logs) + <code>~/Library/Configuration/swarmfile</code> (config) + Keychain items (service <code>com.swarmfile.engine</code>)
Linux	package files, per-user systemd units	<code>~/Library/Caches/swarmfile</code> + <code>~/Library/Configuration/swarmfile</code> ; <code>/etc/swarmfile</code> (purge only); keyring items (service <code>com.swarmfile.engine</code>) where a secret service runs

The cache directory is also where **log files** live, so if you're removing the machine as part of a support case, copy that directory aside first.

Clean reinstall

To start completely fresh on the same machine:

- Follow the uninstall steps above, including deleting the per-user data.
- Install the current package from swarmfile.com/downloads — or, on macOS, make sure it's the `.pkg`: the `.dmg` doesn't set up the engine service or the CLI on `PATH`.
- Sign in again and re-open the project. The mount comes back where the Drive Letter / mount-point settings say; nothing on the hub is affected by the reinstall.

Where to go next

- [Release Channels & Updates](#) — how updates work and what a repair does.
- [Supported Platforms & System Requirements](#) — installer names and per-platform layout.

- [Deploying to Your Team](#) — removing or re-provisioning machines at fleet scale.

Organizations, Projects & Members

An organization is the top-level container for your team: it owns projects, members, billing, and identity configuration. Everything in Swarmfile lives inside exactly one org.

Creating an organization

There's no separate "create your organization" step. When you sign up at `/start`, you register and an org is created for you in the same flow. Choosing **Starter or Pro** goes straight into Stripe Checkout to pick and pay for a plan; choosing **Free** creates the org without any card or checkout. The org gets a slug auto-derived from what you typed at signup; if that slug is already taken, it's auto-resolved to something unique so signup never blocks on a collision.

One account can create up to **25 organizations**; past that, creating another answers with a support pointer. The limit is an abuse bound (a real buyer has a handful of orgs), not a plan feature — contact support if a legitimate setup needs more.

The org slug matters beyond cosmetics — it's part of every public share link (`/s/:orgSlug/:token`).

Renaming your org

As owner, you can rename the org's display name or its slug from Settings at any time.

Renaming the name is cosmetic. Renaming the slug is not: it invalidates every outstanding public share link tied to the old slug, and there's no redirect kept from old to new. Anyone with an old share link loses access silently. Because of that, changing the slug requires retyping the current slug as confirmation before it takes effect — treat it as a real, disruptive action, not a settings tweak.

Projects

A project is a mounted filesystem within an org. Owners and members with sufficient permissions can:

- **Create** a new project
- **Rename** a project
- **Archive** a project (read-only, unmounts for day-to-day use, reversible)
- **Delete** a project (removed from active use straight away; its files, history, git-LFS objects, and stored data are permanently purged 30 days after deletion, and there is no self-serve undelete. Recovery is not guaranteed — [contact us](#) before the purge window closes to discuss it. The 30-day window is fixed; a shorter per-user trash-retention setting does not shorten it. A very large entry tree is removed through a durable background job rather than the immediate cascade, so a big delete can take a while to finish; the CLI waits for it by default and `op status` shows its progress.) Deletion writes durable audit rows that outlive the purge, so it remains auditable after the data is gone.

A project can also be cloned and pushed with git once an owner turns on git access (`swarmfile git enable`, the web project menu, or the Desktop App's **Project settings** → **Git access**) — see [Cloning a Project with Git](#).

Archived projects are read-only

Archiving freezes a project's content until someone unarchives it. Every change is refused with `403 project_archived`, whether it comes from the desktop drive, the web, the CLI, an API key, git-LFS, or a share link:

- saving, uploading, creating, renaming, moving, or deleting files
- commits and changelists, branches (create, merge, archive, protection rules), tags, and merge requests
- publishing releases, RFIs, labels, and comments (adding or resolving, including through a share link)
- changing access: granting permissions, switching the project's ACL mode, or creating new share links (revoking existing access still works)
- taking a lock: file locks, offline checkouts, folder reservations, byte-range locks, and `git lfs lock`

What keeps working:

- reading, browsing, downloading, and history
- **unarchiving** the project
- taking a public release down, revoking access (removing an existing grant or revoking a share link), and releasing locks (including checking a file back in and `git lfs unlock`)
- canceling an upload that was still in progress when the project was archived, so nobody is left waiting for a file that will never arrive

If someone still has the drive mounted, a save to an existing file is kept beside it as a separate copy instead of being lost, and a brand-new file's content stays on their machine with a message explaining why it wasn't uploaded.

Each project runs in one of two modes:

- **Open mode** — the default for every new project; more permissive, suited to small teams or early-stage projects
- **Protected mode** — enforces explicit ACL grants on folders and files; opt-in at creation or by switching the project later

Mode selection and the underlying ACL model are covered in [Permissions](#) — this page only covers project lifecycle, not access control.

Project configuration and templates

One options manifest is the single source for every create-time control — encryption tier, commit mode, retention, visibility, ACL mode, data plane, project type, default-branch protection, the create-only default branch name, **Ensure Git Compatible**, data residency, and the per-project external-collaborator policy. The web dashboard, the Desktop App, and the CLI all render the same set from it, so a choice looks and behaves the same wherever you create the project.

Templates give a new project a starting shape: **Blank**, **Media & post**, **Software (git repository)**, **Revit worksharing**, or **Geospatial**. Where a template defines one, creation seeds a starter ignore file and a folder skeleton; an explicit create option always wins over the template's value.

After creation, the project's settings show each effective value and where it came from. `swarmfile project show` prints the same view for scripts, and `swarmfile project set` changes the settings it can — see the [CLI reference](#) and [Getting Started](#). The billing page's three-state entitlements matrix labels each setting **included**, **upgrade**, or **not built**, so a plan gate is visible before you hit it.

Switching the active organization or project

An engine can host several mounts at once — one project (and one branch) per mount — so a single machine can work in more than one project without a second install. The common case is still one active workspace per engine, and switching it is a hot, in-place operation where possible:

```
swarmfile workspace status          # which org/project this engine is on
swarmfile workspace list            # every org you can reach, plus the current org's projects
swarmfile workspace switch --org <org-id>    # adopt that org's default project
swarmfile workspace switch --org <org-id> --project <project-id>
```

To run more than one project side by side instead of switching, open an additional mount — `swarmfile mounts open <path-or-letter> --label "..."` `--project <project-id>` (same org only) — and target commands at it with the global `--mount <id>` flag; `swarmfile mounts list` shows every open mount and which one is `current`. The per-agent version of this setup is documented in [Coding Agents](#).

Switching is **not** a re-authentication: your session is per-user, not per-org, and already covers every org you belong to. Both same-org and org switches normally happen in place (no engine restart), and a target project whose key has not been released yet is waited for rather than turned into a restart. The engine drains and restarts only when the hot path cannot complete — typically when there is no live mount to re-scope — so the drive may then briefly unmount and return. A hot switch is workspace-wide for the project you are leaving: every mount of that project moves to the new project's active branch (a per-mount branch pin doesn't survive a project change), so open drives follow the switch. A mount on a *different* project stays where it is — and because such a drive can't follow the engine to another **org**, an org switch is refused while one is open (`cross_project_mount_open`). The Desktop App names both before you confirm: the drives that will follow and any that must be closed first. A switch is refused while a changelist is open (submit, cancel, or `--park` it first). A switch doesn't wait for uploads: saves still uploading in the project you're leaving keep uploading in the background, and the Desktop App and `swarmfile status` show them as saving in other projects until they land (see [Reading the sync status](#)). While a switch is running, opening or saving a file on the drive can pause for up to a minute until it finishes, and the Desktop App shows that the drive is switching. A file that was already open in an application before the switch still belongs to the project it came from: its later saves are refused rather than written into the new project, the Desktop App tells you straight away, and when the file closes its unsaved changes are kept on disk and the notification says where. If the application hadn't changed the file before the switch, its first change is refused before anything is written, so Swarmfile has nothing to keep: the changes are still in the application, and the notification says so. Save them under a new name, or switch back and save again. The one exception is an org behind its own SSO — reaching it goes through that org's identity provider, not the built-in one (see [Identity](#)).

A freshly-provisioned machine auto-adopts an org and project on first sign-in, so there's often nothing to switch — see [Deploying to Your Team](#) for how that default is chosen.

Inviting members

Members are invited by email. An invite generates a link; the invitee accepts it to join the org, landing with the coarse `member` org role (below `admin` and `owner` — see the next section for what that hierarchy gates). Don't look for a role picker on the invite screen; there isn't one — you can't invite someone directly as an admin, only promote them afterward. On a plan with ACLs (Pro and above) the invite form can also pre-grant whole-project **Read** or **Write** access, listed on the pending invitation and applied when it is accepted; an invite sent without picks adds the org membership only, and project access is granted afterward under that project's **Permissions** tab. The plan bounds who can be invited: **Free is a solo tier (one seat)** — the Team tab shows a one-seat notice with an upgrade link instead of the invite form — and a Starter/Pro **trial** is capped at two seats until it converts — or until an owner ends it early from the Billing tab, which the Team tab's seat-limit notice links to (admins see an "ask an owner" note instead). A signup that picks more than two seats skips the trial and starts paid. See [Billing & Plans](#). Paid Starter, Pro, and Enterprise orgs have no seat ceiling; each seat simply bills at the plan's per-seat rate. What a member can see and do *within a project* is governed by the ACL system described in [Permissions](#); the org role is a separate, coarser layer that gates org-wide admin surfaces (billing, permissions administration, quarantine) rather than per-file access.

If your org configures SSO with an auto-provisioning group (see [Identity](#)), anyone in that IdP group is added to the org automatically on their first sign-in — no manual email invite needed. SCIM on its own doesn't do this: it syncs your directory into Swarmfile, but org access still requires either that SSO group match or a manual invite. On a public project with access requests enabled, an approved **membership** request arrives through this same invitation flow — the requester accepts the standard email invite to join, and on a protected project the invitation carries the project grant they need to see it ([Public access requests](#)).

Removing members

An owner or admin removes a member from the **Team** tab (only an owner can remove another owner). The hub stops serving them the org's file content and keys within a few seconds. The web dashboard drops the org for them, and their desktop drive stops serving the org's files, including files it had already cached. A computer that's offline at the time does this when it next reaches the hub. A link to a single block that was issued just before the removal can keep working for up to two minutes until it expires.

Their engine also clears its local cache of the project (file listings and cached content) and discards its copy of the project's content key. The one exception is changes they saved that hadn't finished uploading: those stay on their computer, held, and upload if their access is restored. Their desktop app tells them so. The project's key itself isn't changed. For an [end-to-end encrypted](#) project, removal also revokes their key wraps, but can't recall key material already on their device. Re-inviting the member restores their access without a reinstall — including the project permissions they held — and their org role comes from the new invitation.

To cut a removed member off everywhere, Swarmfile visits each of the org's projects in turn, closing their open connections and revoking their keys there; in a large org this can take a few minutes. The same happens when a guest's access is revoked, your directory deprovisions someone, a device key is revoked, or a group's membership changes. Owners and admins can follow it in **Settings → Access revocations**: what's still in progress, any project Swarmfile couldn't reach (with the reason), and what each finished sweep did — including how many devices confirmed they deleted their cached copy, and a **Retry** for a project whose sweep stopped.

Removal also changes the org's peer-to-peer key, the secret your team's computers use to connect directly to each other, so the removed member's computer can no longer join in. Everyone else's desktop app picks up the new key on its own within a couple of minutes, with no restart and no interruption to the drive. If any computer seen in the last seven days is running a client that can't switch keys without restarting, the change waits until those computers update. Owners and admins see a notice with the number of computers in **Settings → Network**, and the org's activity feed records both the postponement and the change when it happens.

Who can see plan and usage standing

Any org member — not just the owner — can see the org's current plan and usage standing: seats, storage, and how close the org is to any limits. This is surfaced in both the Desktop App's plan panel and the web dashboard.

This is a deliberate choice, not an oversight. Owner-only areas of the dashboard — billing and usage administration — stay restricted to owners. Permissions administration and quarantine are open to admins as well as owners. The audit log itself is open to any member (ACL-filtered — you see events on entries you can read); only its CSV export, plus quarantine and org-policy events within the feed, stay owner-only. But plan and usage *visibility* is open to everyone on the org, so a member isn't left guessing why an upload is being throttled or what tier the team is on.

Usage standing is read-only for members; the **request limits** themselves — the per-org and per-user request ceilings — are configurable from the web dashboard's **Settings → Organization** (owners only), the Desktop App's **Settings → Developer** panel (owners and admins), or `swarmfile admin rate-limits` (owners and admins). `swarmfile admin rate-limits usage` shows the current minute's usage against them, and a raise can be time-boxed (`--for 6h`). A single API key can carry its own ceiling (`swarmfile api-key rate-limit`), which is how a render-farm or agent fleet gets more headroom than one person without raising everyone's.

Seat count itself isn't something you set manually — it's derived from live org membership and kept in sync automatically. Details are in [Billing & Plans](#).

Identity

Every org needs a way to authenticate its members. Swarmfile gives you one out of the box and lets you layer on enterprise identity — SSO, SCIM, and on-prem LDAP sync — as your org grows.

Built-in identity provider

Every org gets Swarmfile's own OpenID Connect identity provider by default, with email/password authentication. There's no setup required — it's active from the moment your org exists, and for most Starter-plan teams it's all you need.

Bring your own SSO (OIDC)

Plan: Pro and above.

If your org already runs an identity provider, you can configure it in Swarmfile instead of — or alongside — the built-in one. Owners set it up under **Settings → Identity**: an OIDC issuer URL plus the audience/client id, validated against your IdP's standard `/.well-known/openid-configuration` discovery document and JWKS. Entra ID and Okta are the two explicitly supported and tested patterns via their own OIDC endpoints. **SAML** SSO is offered on Enterprise through an adapter that sits between your SAML IdP and Swarmfile's sign-in, rather than as a self-serve option on this screen; if your IdP only exposes a SAML app, [talk to us](#) and we'll set it up with you.

SSO configuration is per-org, and each org's config is hard-bound: when a member signs in, the token they present is validated against the JWKS for that org's specific configured issuer. One org's SSO setup can't be used to forge identity into another org, even if both orgs use the same upstream IdP vendor.

If an external IdP config is ever misconfigured badly enough to lock everyone in the org out, the owner has a break-glass path back to the built-in Swarmfile IdP. This is an owner-only safety net — don't rely on it as a routine way to bypass SSO.

Desktop app: connecting the desktop app works from an external-IdP (SSO) sign-in too. Swarmfile records the external principal — scoped to that issuer and subject — the first time it sees a verified sign-in, and the device session it then mints is marked as externally originated, so your org's authentication policy exempts it exactly like the SSO browser session it came from. Sign in through SSO in the browser, then use **Connect app** as usual. A subject that collides with an existing Swarmfile account is refused with `identity_conflict` rather than being merged.

By default, SSO only authenticates people who are already members — someone who signs in successfully via your IdP but isn't yet in the org still needs a manual invite. To skip that step, set an auto-provisioning group in the org's SSO configuration: anyone whose token carries that group in its `groups` claim is added to the org automatically as a member the first time they sign in. Leave it unset and nothing changes from the default invite-based flow.

If your org's authentication policy requires MFA, you can also give the SSO config an **MFA authentication context**: Swarmfile sends that value as `acr_values` on the sign-in redirect so your IdP can ask for an MFA-strength authentication. Ask your IdP which value it accepts (for example `urn:mace:incommon:iap:silver`); leave it blank to send none — an IdP that validates `acr` strictly would reject a value it doesn't know, and enforcement at org access is unchanged either way.

SCIM provisioning

Plan: Pro and above.

SCIM 2.0 lets your identity provider push directory sync — creating and disabling principals, syncing group membership — into Swarmfile directly. SCIM by itself doesn't grant org access: a synced user still needs either an SSO auto-provisioning group match (above) or a manual email invite to actually join the org.

Owners manage it under **Settings → Directory Sync**: mint a bearer token for the integration (shown once — copy it into your IdP immediately), revoke one, and watch the recent sync log to confirm users and groups are actually arriving. Point your IdP at the hub's SCIM base URL for your org, `https://<hub>/orgs/<org-id>/scim/v2`; [contact us](#) if you need the exact URL for your deployment. Requests authenticate with the bearer token.

Entra and Okta each use slightly different PATCH conventions for SCIM group and user updates. Swarmfile handles both dialects, so you can point either one at the same endpoint without translation on your end.

On-prem LDAP sync

Plan: Pro and above.

If your org's directory lives on-prem — LDAP or Active Directory rather than a cloud IdP — you run a sync agent on your own infrastructure (for example, as a Docker container pointed at your AD) instead of connecting a hosted IdP. The agent syncs your directory into Swarmfile's principal directory on your schedule, without your directory ever needing to be reachable from outside your network. Like SCIM, this syncs identity — it doesn't by itself grant org access; members still need an SSO auto-provisioning group match or a manual invite.

Authentication policy (require MFA and/or a verified email)

Plan: Starter and above.

Individual members can always protect their own account with an authenticator app (TOTP), under **Settings → Account → Security**. An org can additionally *require* that protection for everyone. Owners configure this under **Settings → Identity → Authentication policy**:

- **Require multi-factor authentication** — members must sign in with an authenticator app. A session signed in before enrollment does not satisfy it; sign in again after enrolling.
- **Require a verified email address** — members must have verified their email before using the org.

You choose when enforcement begins: immediately, or after a grace window (7 days by default; the dashboard offers up to 30 days, and the API accepts up to 3650 days or an explicit enforcement date up to ten years out). During the window everyone keeps working; the web dashboard, the desktop tray, and `swarmfile status` all nudge members to enroll. Once enforcement begins, any org-scoped request from a session — or a personal access token, which snapshots its session's factors

at mint time — that doesn't meet the policy is refused with one clear message and a link to Account settings. Members signing in with a Swarmfile password at the org's login page (`/login/<slug>`) also see the requirement stated up front, so nobody starts a sign-in without their authenticator; members signing in through the org's own SSO provider are exempt from the policy, so no hint is shown on the SSO button.

Members who would fail also get an email: when the policy is enabled, again in the last 24 hours before enforcement begins, and once it takes effect. These are transactional security notices tied to the org's policy, not marketing — they aren't subject to notification quiet hours or unsubscribe settings.

The refusal is never a logout: **enrolling and verifying are never gated by the policy**, and the owner can turn it off at any time from the same Identity page. That off-switch is a genuine break-glass — it stays reachable even for the owner's own non-compliant session.

One gate applies to turning it on: **the owner must have enrolled an authenticator app themselves before requiring MFA for everyone** — the same rule the org enforces, applied to the person setting it. (Signing in through an external IdP exempts you, since your IdP owns that factor.) Once a policy exists, the Identity card also shows an enrollment summary: how many members are enrolled, and which owners and admins aren't yet. **Download CSV** beside it exports the full roster — email, name, role, and enrollment status — for an IT review.

If a member loses both their authenticator app and their recovery codes, an operator can remove MFA for them after verifying who they are — the same effect as the member's own **Turn off** under Account settings: every session is signed out, and the change is recorded in the operator audit trail.

Who's covered:

- **Members and owners, including the owner who enabled it.** Owners are held to their own policy; only the off-switch is exempt.
- **Guests are exempt** — a guest's access is limited to the folder they were invited to.
- **Project API keys are exempt** — they're org-provisioned machine identities, governed by key scoping and revocation.
- **Orgs that sign in through their own IdP are exempt from both flags** (see [Bring your own SSO](#)) — that IdP owns factors and email assertions.

Downgrades follow the same freeze rule as SSO: an enforced policy keeps working if the org drops below Starter, and can always be turned off, but a new requirement can't be turned on until the org is back on Starter or above.

Personal access tokens

Interactive sign-in is the right thing for a person at a keyboard. It's the wrong thing for a script, a scheduled job, or a CLI session that needs to keep working without a browser in the loop. A **personal access token** (PAT) is a long-lived bearer credential for exactly those cases — it stands in for *you*, non-interactively.

The key property is that a PAT carries no permissions of its own. It resolves to your identity, and every access check then runs against **your own current role**, exactly as if you'd signed in normally. A PAT can never do more than you can do right now, and it stops working the moment your access does: change your role and the token follows it; get removed from the org and the token is invalidated along with your membership. There's nothing separate to clean up.

Because it acts as you and dies with your access, minting one is **self-serve for any member** — it isn't owner- or admin-gated. Create and revoke tokens under **Settings → Access Tokens** in the web dashboard. Each token is shown in full exactly once, at creation; after that only its name and prefix are visible. You can hold up to **20 active tokens per member, per org** at a time — the cap is counted against your own tokens in that org, so one member's tokens never consume another's allowance. Revoke one to make room if you hit the cap. Tokens can be given an expiry at creation or left non-expiring for unattended automation; the credential begins with an `sf_pat_` prefix and is sent as a normal bearer token.

If your org **mandates external SSO**, PATs are subject to that too — there's no bypass. A token that would authenticate as your built-in identity is rejected in an SSO-enforced org just as an interactive built-in sign-in would be. A PAT is a convenience for your own access, never a way around org policy.

If your org enforces an [authentication policy](#), the same no-bypass rule applies, with one deliberate nuance: a PAT **snapshots the factors and email assertion of the session that minted it**. A token minted from a compliant session keeps working; a token minted before you enrolled (or before your email was verified) stops satisfying the policy when enforcement begins. Re-sign in with your second factor and mint a fresh token — the Access Tokens page flags any token that won't satisfy the org's current policy.

Security guidance. A PAT is a live credential — treat it like a password. Give each one a descriptive name so you can tell them apart, scope them to a single machine or job rather than reusing one everywhere, set an expiry where the workflow allows, and revoke a token the instant it's no longer needed or might have leaked. Revoking a token from the list ends its access; token material can only ever be created from an interactive sign-in, so revoking every token you can see genuinely ends non-interactive access under your identity.

PATs vs. project API keys

A PAT is not the same thing as a **project-scoped API key**, and the difference matters:

- A **PAT** is *you*, self-served, org-wide within your own access, and tied to your membership.
- An **API key** is a *separate service identity* scoped to a single project, minted by an admin or owner rather than by the member who runs the automation. It has no role of its own beyond that project and doesn't come and go with any one person's membership.
- **Keys are a capped plan resource; PATs aren't.** Each plan includes a fixed headless-key allowance (Free 1; Starter 2 per seat, min 5; Pro 5 per seat, min 10) and extra capacity comes in 10-key packs (\$20/month for 10, either plan); minting past the allowance is refused with `keys_cap` until a key is revoked or a pack is added. Use PATs for a person's own automation and keys for the org's machines — one key per machine or fleet is the recommended shape, for revocation radius and its own rate-limit bucket (see [Billing & Plans](#)).

Revocation takes effect within seconds. Both credentials are checked on every request; revoking a PAT or an API key (or deleting its project) ends its access immediately on the server that handles the revoke, and everywhere else within a few seconds. Revoking a credential also rides Swarmfile's durable revocation sweep, which closes its live event connections — a project API key's own; for a PAT, every live connection for that member, because a connection doesn't name the token that opened it (engines reconnect within seconds with any other valid credential). An org can cap how long one connection may live at all — 5 minutes to 24 hours under **Settings → Organization** — so a connection can never outlast a missed sweep by more than that cap. A changed per-key rate limit applies on the same schedule.

Reach for a PAT when a job should run as a specific human; reach for a project API key when a CI pipeline or service should have its own durable identity independent of any individual. API keys are covered in [Runner as CI](#) and the [CLI reference](#).

Plan gating and what happens on downgrade

SSO, SCIM, and LDAP sync are all Pro+ features, and the restriction is enforced server-side — not just hidden in the UI. Calling an SSO, SCIM, or LDAP endpoint on a Starter-plan org returns a real `402` rejection, not a silently-ignored request.

Downgrades don't retroactively break things. If an org drops below Pro while SSO or SCIM is already configured, the existing configuration keeps working — turning it off automatically would strand people mid-login or silently break provisioning. The gate only applies going forward: on a Starter plan, you can't set up something new, but nothing already running gets switched off underneath you.

For plan comparisons and pricing, see [Billing & Plans](#). For what an authenticated member can actually access once signed in, see [Permissions](#). Setting up SSO with an auto-provisioning group before you push installers to a fleet removes the separate invite step for newly provisioned employees — see [Deploying to Your Team](#).

LDAP Directory Sync

If your users live in on-premises **Active Directory**, this is how their accounts and groups reach Swarmfile. It is AD-specific — the agent reads attributes like `objectGUID`, `objectSid`, `sAMAccountName`, and `uSNChanged`, and generic RFC-2307 directories (`entryUUID` / `uid`) aren't supported today. Swarmfile's control plane can't open raw LDAP connections, so a small **customer-hosted agent** runs inside your network, reads AD, and pushes normalized changes to Swarmfile over HTTPS. Nothing needs inbound access from Swarmfile to your directory.

If your identity provider is cloud-based (Entra ID, Okta), you don't need this agent — those push SCIM directly to Swarmfile. See [Identity](#) for SSO and SCIM setup.

What you need

- A **Pro plan or above** (directory sync is a Pro entitlement).
- An **org owner** to mint the sync token under **Settings → Directory Sync**.
- A **read-only service account** in AD with permission to read users and groups. The agent never writes to your directory.
- A host inside your network to run the agent — Docker is the supported path — with line-of-sight to a domain controller and outbound HTTPS to your hub.

How it behaves

- **Delta sync every 5 minutes** by default (`SYNC_INTERVAL_SECONDS`): scans only entries whose `uSNChanged` is newer than the persisted cursor. Cheap, catches user joiner/mover/leaver churn. Group membership rides along with the full sync, not the delta.
- **Full sync daily** by default (`SYNC_FULL_INTERVAL_SECONDS`): rescans users **and groups**. This is how **deletions** propagate — delta scans can't see a deleted object, so any principal missing from a full scan is disabled (`user_disable` / `group_disable`).
- **Crash-safe**: the cursor is written atomically and a partial batch leaves it unchanged, so the next cycle replays the same window; ingest is idempotent on (`source`, `externalId`), so replays don't duplicate anyone.
- **Disables, not deletes**: an account that disappears from AD is disabled in Swarmfile, preserving its history and audit trail.
- **Server limits**: the hub accepts at most 500 operations per push (batches are clamped to 1–500; default 100) and 1,000 group members per group. The sync token carries its own request-rate ceiling.

Set it up

1. **Mint the sync token.** As an owner: **Settings → Directory Sync → New sync token**. Copy the plaintext (`swarmfile_scim_...`) — it is shown once, and it authorizes principal changes for the org, so treat it like an API key.
2. **Create the AD service account** with read access to the users and groups OUs you'll sync.
3. **Run the agent.** Production-safe Docker Swarm example with secrets:

```
printf 'changeme' | docker secret create ldap_bind_password -
printf 'swarmfile_scim_XXXXXXX' | docker secret create swarmfile_sync_token -

docker service create --name swarmfile-sync \
  --secret ldap_bind_password \
  --secret swarmfile_sync_token \
  --mount type=volume,source=swarmfile-sync-cursor,target=/var/lib/swarmfile-sync \
  -e SWARMFILE_HUB_URL=https://hub.swarmfile.com \
  -e SWARMFILE_ORG_ID=<your-org-uuid> \
  -e SWARMFILE_SYNC_TOKEN_FILE=/run/secrets/swarmfile_sync_token \
  -e LDAP_URL=ldaps://dc01.acme.corp:636 \
  -e LDAP_BIND_DN='CN=svc_swarmfile,OU=ServiceAccounts,DC=acme,DC=corp' \
  -e LDAP_BIND_PASSWORD_FILE=/run/secrets/ldap_bind_password \
  -e LDAP_USERS_BASE_DN='OU=Users,DC=acme,DC=corp' \
  -e LDAP_GROUPS_BASE_DN='OU=Groups,DC=acme,DC=corp' \
  swarmfile/sync-agent:latest
```

Seven variables also accept a `<NAME>_FILE` variant pointing at a file with the value — `SWARMFILE_HUB_URL`, `SWARMFILE_ORG_ID`, `SWARMFILE_SYNC_TOKEN`, `LDAP_URL`, `LDAP_USERS_BASE_DN`, `LDAP_BIND_DN`, and `LDAP_BIND_PASSWORD` — which is the preferred way to keep secrets out of `docker inspect`. The rest (group base DN, filters, cadences, batch size, TLS/timeout, log level) are read from the environment only. If you bind-mount a host directory for the cursor instead of a named volume, `chown 10001:10001` it first — the container runs as UID 10001.

4. **Confirm in the dashboard.** **Settings → Directory Sync** shows the recent sync log; the agent's own logs print a summary per cycle (`users=...` `groups=...` `disables=...` `pushed=...` `failures=...`).

Configuration reference

Variable	Required	Default	Notes
SWARFFILE_HUB_URL	Yes	—	Hub base URL, no trailing slash (production <code>https://hub.swarmfile.com</code> ; use your organization's hub host if different); <code>_FILE</code> variant supported
SWARFFILE_ORG_ID	Yes	—	UUID of the org this agent syncs into; <code>_FILE</code> variant supported
SWARFFILE_SYNC_TOKEN	Yes	—	Bearer token from Settings → Directory Sync (<code>swarmfile_scim_...</code>); <code>_FILE</code> variant supported
LDAP_URL	Yes	—	<code>ldap://host:389</code> or <code>ldaps://host:636</code> ; <code>_FILE</code> variant supported
LDAP_BIND_DN	Yes	—	Service-account DN; <code>_FILE</code> variant supported
LDAP_BIND_PASSWORD	Yes	—	Service-account password; <code>_FILE</code> variant supported
LDAP_USERS_BASE_DN	Yes	—	Search base for users; <code>_FILE</code> variant supported
LDAP_GROUPS_BASE_DN	No	users base	Search base for groups
LDAP_USER_FILTER	No	<code>(&(objectClass=user)!(objectClass=computer))</code>	LDAP filter for users
LDAP_GROUP_FILTER	No	<code>(objectClass=group)</code>	LDAP filter for groups
SYNC_INTERVAL_SECONDS	No	300	Delta cadence
SYNC_FULL_INTERVAL_SECONDS	No	86400	Full cadence (governs deletion detection)
SYNC_BATCH_SIZE	No	100	Ops per push request (server caps at 500)
SYNC_CURSOR_PATH	No	<code>/var/lib/swarmfile-sync/cursor.json</code>	Persist this across restarts
LDAP_TLS_REJECT_UNAUTHORIZED	No	true	Only false for self-signed labs
LDAP_TIMEOUT_MS	No	30000	Connect/search timeout
LOG_LEVEL	No	info	debug / info / warn / error

How AD attributes map

AD attribute	Swarmfile field	Notes
objectGUID	externalId	Canonical mixed-endian GUID, matching PowerShell
objectSid	adSid	Used for Windows DACL projection
displayName → cn → sAMAccountName	Display name	First non-empty wins
mail	Email	Optional
userPrincipalName → sAMAccountName	UPN	First non-empty wins
userAccountControl bit 2	Active (negated)	The ACCOUNTDISABLE flag

For OIDC SSO, the agent pre-links each user's `objectGUID` as the JWT `sub` — the default behavior of AD FS and Entra Connect. If your IdP uses a different `sub` claim, rebind users manually in the dashboard.

Troubleshooting

Symptom	Likely cause
config: required env var X is missing	A required variable isn't set; the agent fails before connecting
bind failed	Wrong bind DN/password, or wrong <code>LDAP_URL</code> ; test with <code>ldapsearch -x -H ldap://host -D '<bind dn>' -W -b '<base>'</code>
/sync/push returned 401	Token wrong, revoked, or minted for another org — mint a fresh one
/sync/push returned 503: tenant provisioning in progress	The org's tenant is still being provisioned; the agent retries next cycle
Users disappeared	The last full sync marked them disabled — check which OU the service account can actually read
Deletions lag	Lower <code>SYNC_FULL_INTERVAL_SECONDS</code> ; note full scans cost O(directory size) per run

Security

- The sync token authorizes **any** principal operation on the org. Store it in a secret manager, never in a committed env file.
- The agent reads AD only; give the service account `Read` on the OUs it needs and nothing more.
- Use `ldaps://` in production — `ldap://` sends bind credentials in the clear. Certificate validation is standard (no skip option beyond the lab flag); there is no custom CA or client-certificate setting today, so your domain controllers need certificates the agent trusts.
- The cursor file holds GUIDs, no secrets; normal filesystem permissions are enough.

Where to go next

- [Identity](#) — SSO, SCIM, and how synced principals join the org.
- [Permissions](#) — grants and protected mode for the synced groups.
- [Security](#) — the broader admin security model.

Permissions

Swarmfile's permission model is built from access control lists (ACLs) applied to folders and files, layered on top of a per-project mode that decides what happens when no grant exists at all.

Read, comment, comment + upload, write, admin

Every grant is one of five levels, and each level can be either an `allow` or a `deny` :

Level	Lets you
<code>read</code>	See and download the contents of the folder or file
<code>comment</code>	Everything <code>read</code> does, plus leave comments and @-mentions on files in the folder
<code>comment_upload</code>	Everything <code>comment</code> does, plus upload new files through the server-side upload route — a guest-reviewer tier, deliberately still below <code>write</code> : it does not grant overwrite, delete, rename, rollback, or restore on existing files
<code>write</code>	Modify contents — create, edit, delete within the folder tree
<code>admin</code>	Manage grants on the folder or file itself, in addition to read/write

Each level implies everything below it, and the levels are ordered exactly as listed above. `comment` and `comment_upload` exist for guest reviewers and external collaborators — see [External Collaborators & Sharing](#) — who need to leave feedback (and, for `comment_upload` , drop off new files) without the ability to touch what's already there.

A grant is set on a specific folder or file and is inheritable down the tree beneath it. You don't have to re-grant access at every nesting level — set it once on a parent folder and everything under it inherits the grant. Depth isn't a tiebreaker: a grant closer to the file doesn't override one set higher up the tree. See below for exactly how multiple grants combine.

A grant applies to the item you name, not to its surroundings. Access is resolved per entry, so granting someone `read` or `write` on a single file lets them open that file (for example, by path or link) without granting the rest of the folder it sits in. The way to it stays walkable, though: a folder on the path to something a person holds a grant on still appears in listings and can be opened, and opening it returns only entries they can read and folders on the path to a grant — everything else stays filtered out. Grant the containing folder when the recipient needs the folder itself — to browse its other contents, or to create entries in it. Inheritance only means something for folders: a file has nothing beneath it to inherit, so a file grant's inheritance flag has no effect, and grants store it off whenever they are written or changed.

How resolution works

When Swarmfile decides whether you can do something to a file, it walks up the ancestor chain from that file to the project root, collecting every grant that applies to you along the way.

Allow grants add up — they don't override each other. If more than one `allow` grant applies along that chain (say, `read` on the project root and `write` on a specific subfolder beneath it), the *highest* level wins, regardless of which grant is closer to the file. A grant lower in the tree can raise your access above what a parent folder gives you; nothing lower in the tree can reduce it. Only a `deny` restricts access below what an `allow` elsewhere in the chain provides.

Deny always wins. If a `deny` shows up anywhere in that ancestor chain, it overrides any `allow` found elsewhere in the chain — no matter which one is "closer" to the file. This is a deliberate, conservative default: a broad `deny` placed high up in a folder tree is a reliable way to lock something down, and nobody can accidentally punch a hole in it with a more specific `allow` lower down.

Where a deny doesn't reach today. Two known limits, so you can plan around them:

- **Read-deny during an identity-service outage.** Grant checks depend on resolving who someone is; during a rare outage of the org's identity/membership service, open-mode reads are served best-effort so a whole project can't lock out, and publishing a public release fails closed in that state. Put folders that must stay restricted in a protected-mode project.
- **Changes reach a mounted drive in two stages.** Granting or revoking access updates the hub immediately (the web dashboard, the CLI, and new mounts see it at once), and a mounted drive is pushed the change, so its cached permission answers refresh on the spot instead of waiting for a remount. A revoke goes further: the revoked folder is removed from that drive's listing and its cached file data is deleted, then confirmed back to the organization as a purge receipt. A newly granted folder is the remaining lag: it can still be missing from that drive's file listing — and writing into it may still need a reopen — until the mount is reopened.

Org owners bypass ACLs entirely. An owner can always get into any project, folder, or file in their org, regardless of what grants exist. ACLs govern members; they don't govern the owner.

Open vs. protected projects

Each project runs in one of two modes, set at the project level:

- **Open mode** — the default for new projects on every plan. Access is permissive: ACLs function as additive grants layered on top of general access, useful for restricting specific sensitive folders without having to explicitly grant everything else.
- **Protected mode** — access is default-deny. Nothing in the project is accessible without an explicit grant. Use this for projects where you want to enumerate exactly who can see what, rather than opt specific things out. Protected is always an explicit choice: the web app and Desktop App offer it next to Open (both default Open), and an API or CLI create that omits the ACL mode gets Open — on every plan, including Pro and Enterprise. Protected mode itself needs Pro or above: creating a project as protected, or switching an existing one **into** protected mode, is refused server-side on Starter. Switching back to open is always allowed.

Switching a project's mode doesn't rewrite existing grants — it changes what happens when no grant matches at all. See [Organizations, Projects & Members](#) for project lifecycle and mode selection.

Access comes from grants, not membership. Joining an org through an invite adds the org role only, so a member with no grant on a protected project sees an empty project — the web files view says so explicitly ("you don't have access to any files here; ask an org admin") rather than pretending the project is empty. On a plan with ACLs (Pro and above), the Team tab's invite form can pre-grant whole-project **Read** or **Write** access alongside the invite, and the Team tab lists every member's whole-project grants next to their name; expanding a member's project there loads that project's folder- and file-level grants on demand, so an admin can inspect the full per-project picture without opening each project (a note appears when a page was capped). Folder- and file-level access is also managed under each project's **Permissions** tab.

Creating at the project root. A protected project's top level has no parent folder for a grant to inherit from, so a new file or folder created directly at the root — or an existing entry moved there — is allowed to an org owner, the project's creator, or a member holding a project-wide `write` (or higher) grant. A folder-scoped grant — even `admin` — does not authorize a root create; work inside an existing folder, or ask an owner to create the top-level folder. Overwriting an existing root file is ordinary `write` access to that file and does not need root authority — it adds no new top-level name. In an open project this is moot: any member who can write creates at the root like anywhere else. The web's Upload / New Folder controls, the desktop drive's root writes, and a `git push` whose commit adds or moves a name at the root follow the same rule.

Publishing a release on a protected project requires project-wide admin access. Org owners and the project's creator always pass, and a project API key passes with a project-wide `write` grant. A release publishes the tagged commit's entire tree to an anonymous URL, and the manifest walk doesn't filter per-entry ACLs, so a member holding only a folder-scoped grant could otherwise publish a tree containing folders they can't read. On an open-mode project, any member who can still write the project as a whole may publish. **Deleting (unpublishing) a release requires the same access as publishing one.** A member made read-only by a deny grant can do neither, and external collaborators can do neither.

You can't publish a release that contains anything you're denied read on. In either mode, if you hold a `read` deny on any file in the tagged tree, or on a folder along its path, the publish is refused with `403`. The deny can be granted to you directly or to one of your groups, and it counts even when it doesn't inherit, because the release's paths would reveal the folder. A deny on a folder the tag doesn't include doesn't block the publish, and neither does a deny below `read` (a `write` deny, say). Ask an org owner or the project's creator to publish instead: a deny never limits what they can read. Unpublishing isn't subject to this check.

A publish fails safe when your permissions can't be checked. If Swarmfile can't look up your organization permissions at that moment, it can't tell whether you hold a deny, so the publish is refused with `503` and the error code `org_unavailable`. It's a temporary refusal, so retry shortly. If your read restrictions cover a very large part of the project, the publish is refused with `403` rather than checked file by file. Neither refusal applies to org owners or the project's creator, and unpublishing is never refused this way.

Creating or deleting a tag requires project-wide write. A tag names the whole tree at a commit, and it's what a release is published from. On an open-mode project that's every member by default. On a protected project you need a project-wide `write` (or higher) grant, since a folder-scoped grant isn't enough; org owners and the project's creator always pass. A project API key needs its own project-wide `write` grant in either mode.

Windows DACL projection

On Windows, Swarmfile projects the effective ACL permissions it's enforcing down onto the actual NTFS-style DACL that Windows and Explorer see on the mounted drive. Permissions aren't just an app-level restriction invisible to the rest of the OS; they show up as real Windows ACL entries on the mount — what you see depends on your own role, by design:

- **Org admins and owners, and anyone with an explicit `admin` grant on a folder**, see the full list of every principal with access to it — exactly what Explorer's Properties → Security tab is built to show.
- **Everyone else** sees just their own effective access level on that folder — read, write, or none — not who else has access. The full roster is deliberately not browsable by ordinary members, the same way it isn't exposed anywhere else in the product: seeing it would let anyone map out an org's whole access structure just by opening folder properties.

Either way, the DACL Windows shows is never stale or invented — it's a live read of what Swarmfile is actually enforcing for the account looking at it.

Branch protection

Folder ACLs govern who can touch *content*; branch protection is a separate governance control over *how changes land*. The day-to-day mechanics — protecting a branch, opening a merge request, the approval flow — are in [Branches & Merging](#). What matters here is that protection is an **org-admin control, not a project ACL**: creating, editing, and removing protection requires an org `admin` (or owner) role, and it's enforced server-side regardless of any grant a member holds inside the project.

There are two ways to protect a branch, and they stack:

- **A per-branch flag** — mark one named branch (say `main`) protected, with its own required-approval count.
- **Glob protection rules** — a pattern like `release/*` protects every branch that matches it, now and in the future, with the rule's own required-approval count.

When both apply to the same branch, the **effective required-approval count is the maximum** of the per-branch flag's count and every matching rule's count (each configurable from 1 to 10). A rule can raise a branch's floor but never lower it below what the branch's own flag already requires.

Protection rules can also require CI checks to pass (`--require-checks-to-pass`), independently of the approval count; the per-branch flag has no checks toggle, so to gate one exact branch on CI use a no-wildcard rule (`swarmfile protection-rule create main --require-checks-to-pass`) — the CLI reference says the same. A branch counts as checks-gated if any matching rule says so — the same OR-based resolution the approval count uses. A CI check that's never reported for the exact commit being merged counts as failing, not skipped, so the gate can't be satisfied by a check nobody's configured to run yet.

A protection rule can additionally carry a list of **auto-reviewers** — a CODEOWNERS-style reviewer list scoped to whichever branches the rule matches. Every merge request opened against a matching branch automatically requests a review from each configured principal, on top of anyone requested by hand. This is bookkeeping, not enforcement: naming someone an auto-reviewer never makes their approval individually required — the required-approval count above is still the only lever for that.

On a protected branch, a direct `merge` is refused, and so is **every direct write** — saving, editing, renaming, moving, deleting, creating, uploading from the web, staging/committing, rolling back, resolving a conflict, auto-resolving one, or reflecting LFS content — with `409 protected_branch`; the change has to go through a **merge request** that collects the required number of approvals before it can land. This is how protection forces review: there's no admin-only "just merge it" bypass on the protected branch itself, and a non-admin also can't archive a protected branch out from under the rule. Folder ACLs still govern who may attempt a write, but a grant cannot override protection.

Limitation. Self-review is blocked *per user* — the person who opened a merge request cannot approve it. That raises the cost of a unilateral change and creates an audit trail, but it is an identity check, not a cryptographic proof that two different humans reviewed the work: a second credential controlled by the same person can still approve. Treat the approval count as a control that raises that cost, not as a guarantee of independent review. Every protection change and every quarantine/ACL action lands in the activity feed — see the audit-log list in [Operations](#).

Plan requirements and downgrades

Folder ACLs and Windows DACL projection are Pro plan and above. They're enforced server-side — the server validates the org's plan entitlement before honoring an ACL operation, so plan gating can't be worked around from the client.

If an org is later downgraded below the tier that supports ACLs, existing grants keep working and nothing is silently opened up. The only thing that changes is:

- **Revoking or narrowing a grant** — always available, on any plan. You can never be stranded unable to fix an over-broad permission just because the org downgraded.
- **Adding a new grant** — blocked on a downgraded plan, until the org upgrades again.

This mirrors the downgrade behavior for identity features — see [Identity](#) — and is intentional: a downgrade should never leave you less secure, only less able to add new access.

For the broader security model — org isolation and storage separation — see [Security](#).

Public access requests

A public project accepts **access requests** from its public page by default (**Project settings** → Public access requests turns them off and on). Any registered Swarmfile user with a verified email — no organization membership — can then ask for one of three levels:

- **Comment access** grants `comment` on the project's root `Contributions` folder — the same [external collaborator](#) tier, without a seat.
- **Contributor access** grants `comment_upload` on that folder: new files only, no overwrite or delete. It is not `write`, and it never grants access outside the folder.
- **Membership** mints the ordinary organization invitation (role `member`), which uses a seat; on a protected project the invitation also carries a whole-project `write` grant so the new member can see it.

Approvals run on the same checks as an owner-minted guest invitation: the external-collaborator ceiling and the [spend cap](#). A refusal on either leaves the request pending with the reason, so buying seats or collaborators and approving again works. Free organizations can grant membership but not comment/contributor access — their public page offers membership only.

The project's owners and admins decide each request from the **Project settings** tab; the queue there and a badge on the Project settings tab show pending work. Guests and contributors an approval creates appear in [External Collaborators & Sharing](#) like any other external collaborator, and revoking them there removes the grant. Approving a guest or contributor request after the project stops being public (or is archived) is refused — use the sharing flow or restore the project instead.

The requester sees the status on the project page and in [My requests](#), can withdraw a pending request, and gets an email on the decision. A declined requester can ask again, with a short message explaining what changed. Requests one person may have pending in an organization at once are capped; when a project is deleted, its requests are removed with it, and old requests age out after 180 days.

Self-Hosted Seed Nodes

Swarmfile is LAN-first: peers on the same network exchange file data directly and only fall back to cloud storage when no local copy is available. A self-hosted seed node extends that further by letting a machine you own — a NAS, an old workstation, a dedicated on-prem box — participate permanently in the swarm as a warm cache tier.

What a seed node is

A seed node is just a machine running Swarmfile with seeding turned on. It joins the peer-to-peer swarm like any other client, but instead of only holding what a given user happens to have opened, it holds a persistent, shared cache that the rest of your office can pull from.

The practical benefit shows up for offices that already have a lot of project data sitting on local infrastructure: rather than every workstation re-fetching the same large files from cloud storage over and over, they fetch from the seed node on the LAN instead. That's faster, and it takes load off your egress and off the cloud fallback path entirely for anything the seed already holds.

A seed keeps two things in sync. It walks the project and fetches **every** block, mirroring the whole project to cloud storage — so the cloud copy is always complete and no editor ever has to be the one who pulls a file from the cloud first. Locally, it keeps a warm cache bounded by `SWARMFILE_CACHE_MAX_BYTES`: once full, the block store evicts the least-recently-used blocks, and an evicted block is simply re-fetched on demand (from a peer or the cloud) the next time someone reads it. Size that budget to the working set you want resident — the template ships 100 GiB; a NAS serving a multi-TB production wants 1 TiB or more. A cache smaller than the project still helps (hot data stays local), but the bigger the cache, the more of the office's reads never touch the WAN at all.

One thing worth being precise about: on the standard managed (encrypted) tier, what the seed node caches is ciphertext — reading it still needs a decryption key fetched live from the hosted hub. A seed node is a warm, on-premises cache that speeds up your office and reduces cloud round-trips, not an independently readable offline copy of your data. See [Data Portability & Offboarding](#) for what that does and doesn't mean if the hosted service is ever unreachable.

Plan requirement

Running a seed node requires the **Pro plan or above**.

It's also mutually exclusive with **cloud-only mode**: a mount can't be both a pure cloud-only client and a seed node at the same time. If a machine is configured for cloud-only operation, you'll need to turn that off before enabling seeding on it.

The hub also gates whether a self-hosted node is allowed to announce itself to the swarm at all, based on your org's plan entitlement. If a seed or NAS node is running under a plan that doesn't include the seeding entitlement, its announces are refused outright — a hard block rather than a warning you might miss in a log.

Enabling seeding

For a machine that's already running the regular Swarmfile client, toggle seeding with:

```
swarmfile seed enable
```

and to turn it back off:

```
swarmfile seed disable
```

Both commands restart the engine to apply the change, so expect a brief interruption on that machine.

For a dedicated, headless machine — a NAS or a box that isn't otherwise being used as anyone's workstation — run the `swarmfile-engine` daemon directly instead of the Desktop App (the Desktop App is just a GUI wrapper around the same engine), then run `swarmfile seed enable` against it exactly as above.

Don't confuse this with the standalone `swarmfile-seed` binary. Despite the name, it isn't a way to run a seed node — it's a one-shot admin tool that uploads a single local file straight into the hub and exits, useful for scripted bulk ingestion. See [swarmfile-seed](#) for its reference.

Enrolling an office cache from the dashboard

For a machine that has no Desktop App — the NAS, the edit-bay Mac mini, the small box next to the switch — set it up from an env file instead. The dashboard generates that file for you:

1. Open **Settings → Office caches** and click **Add an office cache**.
2. Name the office (a stable, lowercase id like `london`) and pick the project the cache serves.
3. The wizard shows the complete `seed.env` once — it contains a freshly minted project-scoped API key. Save it on the cache machine as `/etc/swarmfile/seed.env` (mode 0600), install the engine, and start the engine service with that env file; `SWARMFILE_SEED_MODE=true` is already in it, so there's no toggle to run.

The generated file carries every value a headless seed needs, including the scope (`SWARMFILE_ORG_ID` and `SWARMFILE_PROJECT_ID` alongside `SWARMFILE_OFFICE_ID`) — a seed whose scope is missing pins nothing. The credential is a project-scoped API key: no interactive login on the box, and no refresh token to rotate away. It is an **organization** credential, not a personal one — it keeps working if the person who created it leaves the team, and it stops only when someone revokes it under **Settings → API keys** (at which point the cache shows offline here and the offline alert below fires).

One cache serves one project; run one cache per office per project. The same page lists every cache with its online/offline state, last announce, and the bytes it served this week; a weekly email report goes to the organization's owners and admins, and a cache that goes silent for a day triggers one alert email per outage. Both can be muted under **Settings → Notifications** ("Weekly office cache report", "Office cache offline for a day").

Installing on Windows

On Windows, run the seed as a scheduled task that loads `seed.env` and supervises the engine. Save the generated `seed.env` somewhere the SYSTEM account can read (e.g. `C:\ProgramData\Swarmfile\seed.env`), then create the wrapper and the task from an elevated PowerShell:

```
$wrapper = "$env:ProgramData\Swarmfile\run-seed.cmd"
@"
@echo off
setlocal
for /f "usebackq eol=# tokens=1,* delims==" %A in ("C:\ProgramData\Swarmfile\seed.env") do set "%A=%B"
:loop
"C:\Program Files\Swarmfile\swarmfile-engine.exe"
timeout /t 5 /nobreak >nul
goto loop
"@ | Set-Content -Path $wrapper -Encoding ASCII

$action = New-ScheduledTaskAction -Execute "$env:SystemRoot\System32\cmd.exe" -Argument "/c ``$wrapper``"
$trigger = New-ScheduledTaskTrigger -AtStartup
$settings = New-ScheduledTaskSettingsSet -AllowStartIfOnBatteries -DontStopIfGoingOnBatteries -ExecutionTimeLimit ([TimeSpan]::Zero)
$principal = New-ScheduledTaskPrincipal -UserId "SYSTEM" -LogonType ServiceAccount -RunLevel Highest
Register-ScheduledTask -TaskName "SwarmfileSeed" -Action $action -Trigger $trigger -Settings $settings -Principal $principal
Start-ScheduledTask -TaskName "SwarmfileSeed"
```

The task starts the engine at boot as SYSTEM — no sign-in needed — and the wrapper restarts it if it exits, so a crash recovers headlessly. To remove the seed, run `Stop-ScheduledTask -TaskName SwarmfileSeed` and `Unregister-ScheduledTask -TaskName SwarmfileSeed -Confirm:$false`, then delete the wrapper; the `seed.env` holds a credential, so remove it yourself when you mean to. After the first announce the cache appears on the Office caches page like any other.

Deliberate LAN shard placement

By default, which peer ends up holding which erasure-coded shard of a file is opportunistic — whoever happened to fetch or cache it holds it. That's fine for most teams, but it means fault tolerance is a byproduct of usage patterns rather than something you can rely on.

If you want guaranteed redundancy across your office instead — so that losing any single machine doesn't put a file's availability at risk — turn on deliberate placement:

```
swarmfile ec-placement enable
```

This rendezvous-hashes each shard across your office roster at a replication factor of 3, so every shard has three specific, deterministic homes on your network instead of landing wherever caching happened to put it. Like `swarmfile seed enable`, this restarts the engine.

Reach for this when you have a real office roster of machines and want predictable, guaranteed fault tolerance — not just "probably cached somewhere."

One thing worth being precise about: a seed node's own behavior doesn't depend on whether deliberate placement is on. A seed node always walks the full project tree and fetches every block, unconditionally, so the cloud copy is guaranteed complete — that's a stronger, whole-project guarantee than deliberate placement's replication-factor-of-3, not something deliberate placement adds to. What stays on the seed's own disk is its warm cache (bounded by `SWARMFILE_CACHE_MAX_BYTES`, see above), not an unbounded mirror: "complete" here means every block has been fetched and mirrored by the seed, with recently used blocks resident for LAN reads. Deliberate placement is what upgrades your *regular* (non-seed) peers from opportunistic to guaranteed coverage; if every machine in your roster is already a seed node, the cloud always has everything and nothing is left to chance for recovery.

This doesn't require a literal office LAN. `office_id` is a grouping label, not a network check — any set of peers that share one `office_id` join the same placement roster together, whether or not they can reach each other on a local network at all. That's the pattern for a cloud-hosted render farm: set a matching `office_id` on each node, and leave `--lan-from-office` off (that flag exists to fold regular peers on a real LAN into the roster when mDNS can't discover them; it's not relevant when there's no LAN). Leaving it off also keeps these nodes correctly classified as WAN traffic rather than wrongly exempted from the local bandwidth throttle — egress is never billed on any plan either way. See [Deployment Topologies](#) for the render-farm case end to end.

For context, files are erasure-coded Reed-Solomon 10+4 by default, and that scheme is adaptive: it exists primarily to mask WAN latency, so once two or more LAN peers are present for a project it can automatically back off from the full redundancy it uses over a pure cloud connection. Deliberate placement overrides that automatic behavior — when you explicitly want guaranteed office-wide redundancy rather than whatever the adaptive default decides, `ec-placement` takes the decision out of the engine's hands. For more on the underlying security and isolation model, see [Security](#).

Inspecting shard placement

To see exactly where a given file's shards actually live, run:

```
swarmfile shards <path>
```

This shows, per shard: which peer holds it, whether it's a data shard or a parity shard, and whether the machine you ran the command on currently holds a copy itself. It's the tool to reach for when you want to confirm deliberate placement actually took effect, or just to understand a file's current redundancy without guessing.

Security

This page is the deep-dive companion to the [Security](#) overview on the marketing site. It covers how encryption, integrity checking, and network isolation actually work, and it's explicit about what's available today versus what's still on the roadmap.

Running a formal vendor security review? See [Security Architecture](#) — the threat model, a who-can-decrypt matrix across the managed and E2E tiers, key-custody specifics, and where each guarantee's limits are, written to hand to a security team.

Encryption

Every **private** project on a paid plan (Starter and above) has its block contents encrypted at rest with AES-256-GCM by default. Each project gets its own key, wrapped by a hub-held key-encryption-key (KEK). Desktop writes are encrypted on your machine before upload; small files uploaded from the web are sent over TLS and encrypted by the hub the moment they arrive — so stored private blocks are always ciphertext. Because the hub holds the KEK for a managed project, Swarmfile can decrypt when it needs to (for example, to generate a preview or thumbnail). So managed encryption protects your data at rest from the storage layer, not from Swarmfile itself — for that stronger guarantee, use the end-to-end tier below.

Two cases are stored unencrypted (plaintext). The **Free plan** is one: its projects are plaintext at rest, the trade-off for no-card, no-cost storage. The other is a **public project** on any plan — anonymous public releases can only be served from unencrypted blocks, so a project has to be plaintext to be public (Free-plan projects are public by definition). Encryption tier and visibility are both fixed at project creation and never change afterward, so a project created on the Free plan or as public stays unencrypted even after the org upgrades; a new private project created post-upgrade gets managed encryption normally. If encryption at rest matters for a given project, create it private on Starter or above. Public **exposure** is still scoped by the publisher: a committed `.swarmfile/publicignore.yml` keeps matched paths — and derived data such as CI logs — out of what non-members can reach, enforced as each release is published (share links are not covered; see [Security Architecture](#)).

One carve-out inside an encrypted project: some git-LFS objects. When a managed project is used as a [git-LFS server](#), objects uploaded directly by the stock `git-lfs` client, and uploads of 64 MiB or more through the Desktop App's built-in LFS agent, are stored **without application-layer encryption**. Smaller uploads through the built-in agent, and every upload through the standalone `swarmfile-lfs` agent, are encrypted like any other block. git-LFS is refused on end-to-end-encrypted projects. Files you save through the mounted drive are unaffected; if a binary must be encrypted at rest, keep it on the drive rather than routing it through LFS. The [git-LFS guide](#) has the full per-path table.

Opt-in end-to-end (customer-managed) encryption

For projects that need a stronger guarantee than a hub-held key, you can opt into end-to-end encryption. Instead of wrapping the project key with a hub-held KEK, each member's copy of the key is wrapped to their own device keypair (X25519). The key never reaches Swarmfile's servers in plaintext at all, and grants happen client-side.

This is a content-only guarantee:

- File and folder **contents** become unreadable to the server.
- Filenames, folder structure, and file sizes stay server-visible — search, ACLs, browsing, and quota accounting keep working normally.
- Stored blocks are named by a hash of their plaintext chunk (so identical content deduplicates). No whole-file hash is kept for E2E content — see [Security Architecture](#) for exactly what that does and doesn't reveal.

The trade-offs are real and worth knowing before you turn it on:

- Server-side preview and thumbnail generation don't work for E2E content, because the server can't decrypt it.
- New members are enrolled into E2E projects automatically once an existing member's client has been online to complete the key wrap — no manual key ceremony. An optional always-on key-granter closes even that narrow window, delivering new-member and new-device keys automatically on a five-minute cadence at worst. See [End-to-End Encryption Setup](#) for the granter setup.

Losing every device with a wrapped copy of a project key would mean losing the data permanently, so E2E projects should be created under an org recovery key. That key is split using Shamir secret-sharing escrow (2-of-3 by default), so no single lost device or credential is catastrophic.

For the step-by-step — setting up the recovery key, creating an E2E project, enrollment, and the recovery ceremony — see [End-to-End Encryption Setup](#).

Content integrity

Every block is BLAKE3-verified on arrival. Corrupt data — from a misbehaving peer, a bad transport, or disk corruption in transit — never reaches the application layer.

Erasure coding

Blocks are protected with adaptive Reed-Solomon 10+4 erasure coding: any 10 of the 14 shards reconstruct the full data, so up to 4 slow or offline peers never block your read. This page keeps that summary brief — for the admin-facing detail on deliberate shard placement across your seed nodes, see [Self-Hosted Seed Nodes](#).

Private P2P swarm

The peer-to-peer transport runs in a private swarm. Joining it requires a pre-shared key and a custom protocol identifier — a random host on the public internet can't dial into your blocks even if it somehow learned a peer's network address. Peers that connect directly do learn each other's network addresses, which is inherent to a direct connection; an org that can't accept that can turn on cloud-only mode (below), which disables peer-to-peer entirely.

Cloud-only (hub-only) data-plane mode

Some enterprises need to block UDP/QUIC entirely, or eliminate peer-IP exposure between employees' machines outright. An org-level policy — or an engine-local `SWARMFILE_HUB_ONLY` toggle — disables peer-to-peer entirely: QUIC never binds, there's no peer discovery, no gossip, and `swarmfile doctor` reports clean. All reads

fall through to HTTPS against our cloud storage instead. This is enforced at startup and re-checked every time a machine switches which project or organization it's working in, so a person who works across multiple orgs never inherits the previous org's policy; a check that can't complete (a hub outage, for instance) fails closed rather than silently allowing peer traffic.

This is admin-gated (owner or admin), and toggling it is itself an audited event. See [Operations](#) for where that shows up in your activity log.

Removing access

Revoking access reaches every affected computer, not just the next sign-in. Remove someone from the org, revoke their access to a folder, or revoke a machine credential, and each affected drive drops the revoked content from its local cache and refuses further opens of it — including while offline, because the refusal is persisted locally the moment the notice arrives. A machine that was offline takes the same step when it reconnects. Organizations can also bound how long a *disconnected* machine may keep serving already-downloaded content at all (default: seven days); past that window the drive stops serving cached bytes until it next reaches the hub. Removed members' own unsaved work stays on their machine and uploads if their access is restored. Owners can see any revocation still pending under **Settings → Access revocations**.

Ransomware detection

The system watches for a suspicious pattern inline, at high frequency: a large number of *distinct entries* overwritten (by content-hash) in a short window, scoped per-(user, machine). It counts distinct entries, not raw history rows, so re-saving the same file repeatedly doesn't inflate the signal.

The response is two-tier, and crossing the first line does **not** block anyone:

- **Alert bar** (a configurable number of distinct entries in the window) — the incident is recorded as alert-only and org owners are notified for review. Nothing is blocked.
- **Block bar** (a higher multiple of the alert bar) — this is what actually hard-quarantines the user, either on a genuinely large sweep or when an already-alerted user keeps going. Only past this bar do writes stop.

Writes are attributed to the machine that made them. A staged commit counts in the overwrite signal by its modifications — the submitting machine rides the submit, and the check runs as the commit lands, so a single 500-entry commit from one machine blocks on the spot (250–499 alerts) — while the server-side re-applies a restore, checkout, revert or merge performs (content the project already held) are excluded, so a legitimate 5,000-file restore does not trip the overwrite bars. A staged commit's creates are creates, never overwrites. A `git push` carries the pushing machine like any other client batch, so push sweeps are counted too. Two accepted exceptions remain: the web app's conflict tools (the resolve modal and the batch auto-resolve) carry no machine and are not counted, so a bulk auto-resolve is invisible to this signal; the separate create-and-delete signature is evaluated inline as a direct delete lands, while a merge that both adds and removes hundreds of entries — applied server-side with no inline check — and an over-cap folder-delete subtree reach it only through the daily scan, so coverage there is best-effort. A delete-only sweep has no signature at all and is not counted by design — deletes alone are how anyone empties a folder. An in-place sweep of hundreds of distinct files from one client can still reach either bar; orgs whose workflows routinely rewrite that many files in place can tune the thresholds.

When a user is quarantined, lock acquire and lock renew both check quarantine state, so they can't keep working even in the middle of holding a lock. New saves stop on that computer: on macOS and Linux they fail with the `quarantined` reason in **Why did my save fail?**, and on Windows the drive stays mounted while saves are accepted locally and held — they simply don't sync until the block is released. Work already queued on that machine is held, not lost — sync resumes on its own as soon as an owner releases the block. Alert-only incidents that no one acts on auto-expire after 7 days; blocks never auto-expire and require an admin or owner to clear. See [Operations](#) for the release/dismiss workflow.

No kernel extension on macOS

The macOS installer ships a kext-free build using FUSE-T, which mounts entirely in user space. There's no kernel extension, no dropping into Recovery Mode to lower security settings, and no system-extension approval prompt on Apple Silicon.

Authentication and membership

- **Sign-in** runs on Swarmfile's own OIDC IdP, and every account can add TOTP multi-factor authentication under Account → Security.
- **Organizations can require MFA and/or a verified email** — see [Authentication policy](#). Enforcement is checked on every org-scoped API request with a grace window and a typed refusal that names the remedy; the owner always keeps an off-switch. A personal access token snapshots its session's factors and email assertion at mint time, so a token can never carry more than the session that created it.
- **SSO**: per-org OIDC (Entra/Okta), SCIM 2.0 provisioning, and an on-prem LDAP sync agent; an org can require that members sign in through its own identity provider, and machine credentials are held to the same binding.
- **Machine credentials**: project-scoped API keys for headless engines and personal access tokens for scripts; both are revocable and audited. API keys are additionally walled to one project and can carry their own budget; a personal access token acts as its owner and can never carry more than the session that minted it.

On the roadmap

- **SOC 2** is on our roadmap; we have not yet begun a formal audit. In the meantime, the team answers security questionnaires directly and walks customers through the architecture on request.
- **SIEM forwarding** — a managed Splunk, Datadog, or S3 stream of audit, access, and quarantine events — is planned. Today, a Pro-plan owner can export the [activity feed](#) to CSV (up to 100k rows per export), and [outbound webhooks](#) push events to an endpoint you control on a periodic schedule, not in real time.
- **Legal hold & retention immutability** — a per-entry legal-hold flag that overrides retention and trash purge — is planned.
- **DLP & egress controls** — watermarking, share-link domain allowlists, and export policy gating — are planned.

If any of these are a hard requirement for you today, talk to us directly — for details on the permission model that is live today, see [Permissions](#).

Security Architecture

This is the deep technical companion to [Security](#) — written to be handed directly to a security team running a vendor review. It states plainly what each layer protects against, who can decrypt what, and where the limits of that protection are. If your review needs something not covered here, [contact us](#) and we'll answer specifics.

The one-paragraph version

On a **paid plan** (Starter and above), **private** project content is encrypted with AES-256-GCM before it is stored — on the machine that writes it for anything saved through the drive, and by the hub for browser uploads to a managed project (the hub holds a managed project's key; an end-to-end project's key never reaches it). Peer-to-peer nodes and our cloud object storage see only ciphertext addressed by a BLAKE3 content hash, with a few documented exceptions: the [git-LFS](#) direct-upload path (see below); server-side **video proxies** and the **whole-file copies** assembled to serve browser downloads (both retained short-term — see "Worth knowing about server-side previews" below); and any **public** project, which is stored unencrypted by design — see [Unencrypted tier](#). Free-plan projects are all public and use that tier because they have no other, and a paying org can still choose it deliberately for a public project that needs the anonymous-download paths (which read unencrypted content only). Because a public project's content is public, its plaintext blocks may also be served to peers on the org's private swarm — peers reach only what the anonymous read routes already serve, peer serving is read-only, and every fetched block is BLAKE3-verified before it is used. A paying org's **private** projects are encrypted — managed by default, E2E opt-in — which is also what makes them eligible for peer-to-peer transfer. The **only** party that can turn ciphertext back into your files is whoever holds the project key — and *you choose who that is*. On the default **managed** tier that party is the Swarmfile hub, so we can give you server-side previews, search, and proxies — managed means *encrypted at rest with hub-held keys*, not zero-knowledge. On the opt-in **end-to-end (E2E)** tier the key never reaches our servers in a form we can open, so we structurally cannot read your content. Both tiers are fully supported; the difference is a deliberate trade you make per project, not a weakness in one of them.

The **Free** plan is the one plan with no at-rest encryption at all: its projects are stored **unencrypted**, with no project key — the trade-off for no-card, no-cost storage. See [Unencrypted tier](#) below. On Starter and above, a **private** project gets managed encryption by default (E2E opt-in); a **public** project is unencrypted-only on every plan, because public releases and archives still read unencrypted content only. The unencrypted tier is therefore Free-plan projects plus a paying org's deliberately-public projects.

Threat model — what each tier defends against

Adversary	Managed tier	E2E tier	Unencrypted tier (Free plan; a paid org's explicit public-project opt-out)
A malicious/curious peer on your swarm	✅ sees only ciphertext + BLAKE3 CIDs	✅ same	❌ sees plaintext content — which is public by design (peers see only what anonymous readers see)
The storage provider or anyone who exfiltrates a storage bucket	✅ ciphertext at rest, no keys co-located	✅ same	❌ plaintext at rest
Someone who subpoenas or compromises the Swarmfile hub	⚠️ hub holds the wrapping key and <i>can</i> decrypt	✅ hub holds only un-openable wrapped blobs	❌ nothing to compromise — already plaintext
A Swarmfile employee with production access	⚠️ same as above — possible, audited, not prevented	✅ prevented by construction	❌ same as above
A non-member of the project (incl. other orgs, other members)	✅ blocked by the block ACL on every hub fetch	✅ blocked by ACL <i>and</i> has no wrapped key	✅ nothing confidential to reach (public by design); hub fetches stay ACL-gated, and swarm peers can only fetch blocks the anonymous read routes already serve
A lost password / lost device	✅ recoverable (hub re-wraps)	⚠️ recoverable only via the org recovery key	✅ no key to lose

Exception — git-LFS objects. The block properties above describe project content written through the drive/engine. A project used as a [git-LFS](#) server is not uniformly encrypted at rest: objects uploaded through an agent's chunked block path are encrypted on a managed project (the desktop agent client-side; the standalone agent's plaintext blocks are encrypted by the hub on write), but two direct-upload paths store the object unencrypted — a stock `git-lfs push`, and the desktop app's built-in transfer agent for objects of 64 MiB or more. That is one of the documented exceptions to "ciphertext at rest", and it is a deliberate trade: LFS is an add-on surface, and encrypting and decrypting multi-GB artifacts on every transfer isn't worth it. A binary that must be encrypted at rest belongs on the mounted drive (managed or E2E), not routed through LFS.

The two ⚠️ rows on the managed tier are the entire reason the E2E tier exists. If "the vendor can technically decrypt our footage" fails your review, the answer is the E2E tier — not a promise about the managed one. The unencrypted tier is a different trade entirely: it is not a weaker version of managed encryption, it is the deliberate absence of at-rest encryption in exchange for a free, no-card tier — access is still gated by the same block ACL every other tier uses, but there is no encryption layer to defend the *storage* boundary at all.

How block encryption works (both tiers)

Encryption of the bytes is identical on both tiers; only *custody of the project key* differs.

- Each project has one 32-byte **project master key**.
- Every content block gets its own key, derived from the project key and bound to the block's content (the CID is the BLAKE3 hash of the *plaintext* chunk), so a key is never reused across blocks.
- The block is sealed with **AES-256-GCM** and keyed in object storage by its CID. (The exact block format is in the [Storage Format](#) spec.)
- On read, the reader re-derives the key, decrypts, and re-verifies the BLAKE3 CID before any byte reaches an application — so a tampered or corrupted block is rejected, not served.

Content-defined chunking means identical content across file revisions shares a CID and deduplicates, and the per-project key scoping means dedup never crosses a project boundary.

Managed tier (default)

The project master key is generated once, **wrapped by a hub-held key-encryption-key (KEK)**, and stored as ciphertext in the org's metadata store. The hub unwraps it only when it needs to act on your behalf — serving a download, generating a thumbnail, or transcoding a preview proxy.

What this protects: your data at rest against the storage layer, a leaked bucket, a lost disk, or a peer on your swarm. What it does **not** protect against: the hub itself, since the hub custodies the KEK. That is the standard posture of essentially every managed B2B file service (it is what lets us offer previews and search) — and it is why, between two members of the same org, the confidentiality boundary is the **block ACL**, not cryptography. Every block fetch is authorized per CID against the project's permission model; encryption is the at-rest guarantee, the ACL is the member-to-member guarantee.

Web uploads (small files sent from the browser over TLS) are encrypted by the hub the instant they arrive, before anything is written to storage — so managed data is always ciphertext at rest, matching what the desktop engine writes client-side.

Worth knowing about server-side previews. Because the hub *can* decrypt managed content, it renders previews on the server: thumbnails, video poster frames, point-cloud images, image/PDF previews, and scrubbable video proxies. The rendering happens in memory. What gets stored depends on the kind of preview:

- **Thumbnails, poster frames and point-cloud preview images are stored encrypted** under the project's key, with the same AES-256-GCM block encryption as your files, and decrypted only to serve a request that already passed the file's read check. The hub can still decrypt them, exactly as it can decrypt the source files. That is what managed encryption means.
- **Still stored unencrypted:** scrubbable **video proxies** (a 720p transcode, kept up to 30 days), the **whole-file copies** assembled to serve browser downloads and previews (kept up to 24 hours), and the short-lived inputs and outputs of a preview render, deleted when the render finishes.

So a compromise of the storage layer could still expose a video proxy or a recently downloaded file, even though the source blocks stay sealed. The **E2E tier** removes this residual entirely: on E2E the hub can't decrypt, so no server-side previews or derived copies exist. If preview confidentiality at rest matters for your content, choose E2E.

Key rotation

Two things can rotate: the KEK and the project key. Rotating the KEK only re-wraps the stored project keys. The keys themselves don't change, so no block needs re-encrypting.

A managed project's key can also be rotated — by the project's owners, its admins, or its creator, from the dashboard's project settings. This adds a new key *generation* and makes it current. Earlier generations are kept, so every block written before the rotation stays readable, and each block records which generation sealed it (see [Storage Format](#)). New content is sealed under the new generation by every client that knows about key generations — the Desktop App, the hub itself, and the git-LFS agent. A rotation doesn't re-encrypt content that's already stored, though, and an old generation is never destroyed. So after a suspected key compromise, content written before the rotation is still protected only by the old key.

An **E2E** project rotates the same way — the generation model is identical — but the hub never holds or serves the key, so a member's own device performs the rotation: it mints a new generation and re-wraps it to every remaining device plus the org recovery key, and the hub only stores the envelopes the device produced. The same caveat applies: existing content is not re-encrypted, so it stays readable under the old generation. Revoking a single device is separate — the hub revokes that device's wrap(s), and the wrapped-key route then answers it a typed `key_access_revoked`.

Every key-encryption-key (KEK) is a **ring**, not a single value: new values are wrapped under the active entry and reads try every entry, so a rotation never makes stored data unreadable. The same discipline covers the managed-tier project-key KEK and every credential the hub stores for a customer.

A rotation is two-step, and the old entry is never dropped in the same step that adds its replacement: stored values are re-sealed under the new entry first, and only then is the old key retired.

Unencrypted tier

The Free plan — no card required — is restricted to a third tier that is not a variant of managed encryption: it has **no project key at all**. Blocks are written and stored exactly as they arrive, no AES-256-GCM sealing, no KEK, nothing to wrap or unwrap. A **public project on any plan** also has to be on this tier, for the same underlying reason: anonymous public serving only works from unencrypted blocks. This is the tier's design, chosen deliberately so a free tier can exist without a per-project key — and, for public projects, because there is nothing to hide from anonymous visitors anyway.

What stays the same as every other tier: the **block ACL**. An unencrypted project's content is not encrypted, but it is not *unauthenticated* either — every hub block fetch is still authorized per CID against the project's permission model, the same gate managed and E2E projects use. Plaintext is peer-served only for **public** projects, and then only with content the anonymous read routes already serve, while a private project's cleartext is never served there. Plaintext removes the storage-at-rest guarantee; it does not remove access control on the hub.

Encryption tier is fixed at project creation and immutable afterward. A project created on The Free plan stays unencrypted even after the org upgrades to a paid plan — upgrading does not retroactively encrypt existing content. A new **private** project created after upgrading gets managed encryption normally, the same as any Starter/Pro project. If a specific project needs encryption at rest, create it on a paid plan; don't rely on a later upgrade to change an existing Free-plan project's guarantee.

The Free plan can only create **public** projects (private projects require Starter or above), so its unencrypted-tier content is also, by definition, intentionally public-visible. On a paid plan a project is private and managed-encrypted by default, and only enters this tier if it is created as **public** — visibility is fixed at creation, never flipped, so this is a create-time choice: the CLI's `swarmfile project create <name> --public` (which pins the unencrypted tier with it) or the create API's `"visibility": "public"` + `"encryptionTier": "plaintext"`. The Desktop App's New Project dialog and the web dashboard both create private projects by default, with a **Publish publicly** card for a project meant to be open (unencrypted projects remain the Free-plan shape there). See [Permissions](#) for how project visibility works; it's a separate setting from, but always paired with the unencrypted tier.

Keeping paths out of the public view

An unencrypted public project's content is public by design, but a publisher chooses *which* content actually goes public. A release reads a `.swarmfile/publicignore.yml` committed at the tag's tree and leaves matched paths out of the public release entirely — no browse entry, raw stream, download, archive, file search or Explore — while members keep seeing everything. The same policy can mark platform-derived data (currently CI-run logs) member-only. The policy file itself is never published, an invalid file fails the publish instead of publishing everything, and a change applies to future releases only, so an already-published release needs a new tag to pick it up. The full syntax and the enforcement points are in [Publishing Releases](#).

Two boundaries worth stating plainly: the filter applies to the public release surface only — **share links are not covered**, so a link still serves what it points at — and an org policy can only *tighten* this exposure, never loosen it. It is a publisher control, not a retroactive takedown: copies already downloaded stay where they are.

End-to-end (E2E) tier (opt-in, per project)

On an E2E project the project key is **never wrapped with a hub-held KEK**. Instead:

- Each member has an **X25519 keypair**, generated and enrolled per device; the private key never leaves the member's machine. (Deriving the same keypair from a password via Argon2id — so a member could recover access on a new device without a per-device re-enrollment step — is on the roadmap, not implemented today: every device holds its own independently-generated keypair.)
- The project key is **wrapped to each member's public key**. Granting a new member = an existing member's client wraps the project key to the newcomer's public key and uploads the (un-openable) wrapped blob. The hub distributes wrapped blobs but never sees the plaintext project key.

This is a **content-only** guarantee, and we say so up front: file and folder *contents* become unreadable to the server, but **filenames, folder structure, and file sizes stay server-visible** so that search, ACLs, browsing, and quota accounting keep working. Encrypting metadata too would break all of those and is a separate, much larger effort — nearly every B2B "E2E" product draws this same line.

To be precise about which metadata: what the server sees is the *entry-level* metadata it keeps in its database — the name, its place in the folder tree, and the file's total size. It does **not** see the file's internal structure. The per-file block **manifest** (the ordered list of content-chunk hashes and their sizes) is itself encrypted alongside the content — the engine runs it through the same block-encryption as the chunks — so on the E2E tier the server holds the total size but the chunk-level layout is opaque to it. (This is why an E2E share is fetched and reassembled entirely in the recipient's browser: the server can't read the manifest to hand out a chunk list, let alone assemble the file.)

One more thing the server does see: **block names**. Each stored block is named by the BLAKE3 hash of its *plaintext* chunk — that is what lets identical content deduplicate. So the server can tell when two stored blocks are the same, and someone who already holds a copy of a file could chunk it the way the engine does and check whether those blocks exist in a project. What the server does **not** keep for E2E content is any whole-file hash: the SHA-256 that git-LFS uses to address a file is never stored for an E2E project, so E2E files can't be matched against published lists of file hashes.

Share links on an E2E project

A public share link has no account and no device keypair, so it can't receive the project key wrapped the member way. Instead the link carries the project key wrapped to a **per-share random "link key"** that lives only in the URL `#fragment` — browsers never send a fragment to a server, so the hub distributes the wrapped blob but never sees the link key or the project key. The recipient's browser unwraps it and decrypts the file locally.

The link carries the **project master key**, not a per-file key — the block cipher derives every block's key from the project key, so there is no narrower key to hand out. Two consequences follow:

- **The share's block endpoint is scoped to the shared file.** The hub serves a share's raw ciphertext blocks through one endpoint, restricted to the shared file's own block set: its manifest plus the chunk hashes the creator's browser enumerated at share time (the hub can't read the E2E manifest to derive them itself).
- **That boundary is enforced at the serving layer.** Because the recipient genuinely holds the project key, the hub's scoping of the link to the shared file's blocks is an access control, not cryptographic isolation. Per-file share keys — which would make that isolation cryptographic — are a planned change. Treat a single-file share as granting access to that file and nothing else: scope links with expiry and revocation, and share narrowly. Sharing is an explicit trust decision.

The trades E2E makes (know these before you turn it on)

- **Server-side previews/proxies are disabled** for E2E content — the server can't decrypt to render them. Image/PDF preview can move to in-browser decryption; server-generated video posters and point-cloud rasters cannot. This is a deliberate, headline trade-off.
- **SSO/SCIM grants authentication, not key access.** A provisioned user can sign in but sees nothing until an existing member's client (or the optional key-granter, below) wraps the project key to them; when no grant was queued at all (a group-added member, a failed fan-out), the member's own engine asks the hub to open one.
- **Recovery is via an org recovery key, or not at all.** Losing every device holding a wrapped copy means losing the data. The enterprise-viable answer is an **org recovery key**: project keys are *also* wrapped to a recovery public key whose private half is escrowed via **Shamir secret sharing (2-of-3 by default)** across your admins. This reintroduces a party that *can* decrypt, so "E2E with recovery" is a point on a spectrum — customers who choose no-recovery accept unrecoverable loss, and we describe it that way at the point of choice. Rotating the recovery key does not re-wrap existing projects: a project stays sealed to the key in force when it was created, recovery accepts shares from any recovery key the org has created, and the previous shares must be kept or that project becomes unrecoverable.

Optional always-on key-granter

For orgs that add many members while their existing members are offline, an opt-in **key-granter** service can hold an escrowed recovery share and fulfill pending grants on a tight cadence. It is off by default and, like the recovery key, is a decrypting party you are choosing to introduce for operational convenience — enrolled explicitly per org, never automatically.

Integrity, availability, and network isolation

- **Content integrity.** Every block is BLAKE3-verified on arrival; corrupt data from a bad peer, transport, or disk never reaches the application layer.
- **Erasur coding.** Blocks are protected with adaptive Reed-Solomon 10+4 — any 10 of 14 shards reconstruct the data, so up to 4 slow or offline peers never block a read.
- **Private P2P swarm.** The peer transport requires a pre-shared key and a custom protocol identifier; a random internet host cannot dial your blocks even if it learns a peer's address. Peers that connect directly do learn each other's network addresses, which is inherent to peer-to-peer transfer; the cloud-only mode below removes that exposure.
- **Cloud-only (hub-only) data-plane mode.** An org policy (or engine-local `SWARF_FILE_HUB_ONLY`) disables P2P entirely — QUIC never binds, no peer discovery, no gossip, no peer-IP exposure — and all reads fall through to HTTPS against cloud storage. Re-verified fail-closed on every switch between projects or orgs, not just at startup, so a policy check that can't complete never leaves P2P silently enabled. Toggling it is itself an audited event.

- **Ransomware detection.** The system watches for the signature of encryption malware — a large number of *distinct entries* overwritten (by content hash) in a short window, scoped per-(user, machine) and counting distinct entries rather than raw history rows. The response is two-tier: crossing the **alert bar** (a configurable number of distinct entries) only records an alert-only incident and notifies org owners — it does *not* block — while crossing the **block bar** (a higher multiple of it) is what hard-quarantines the user, halting further writes even mid-lock. Writes are attributed to the machine that made them: a staged commit carries its submitting machine, its modifications count in the overwrite signal (the check runs as the commit lands, so a single 500-entry commit from one machine blocks immediately), and its creates are creates, never overwrites — while the server-side re-applies a restore, checkout, revert or merge performs (content the project already held) are excluded, so a legitimate 5,000-file restore does not trip those bars. A `git push` carries the pushing machine like any other client batch, so push sweeps count as well. Two accepted exceptions remain: web-app conflict tools (including the batch auto-resolve) carry no machine and are not counted; the separate create-and-delete signature is evaluated inline as a direct delete lands, while a merge that both adds and removes hundreds of entries — applied server-side with no inline check — and an over-cap folder-delete subtree reach it only through the daily scan, so that coverage is best-effort. A delete-only sweep has no signature at all — deletes alone are how anyone empties a folder, and the conjunction is what distinguishes it from encryption-to-new-names. An in-place sweep of hundreds of distinct files from one client can still reach either bar; orgs that routinely rewrite that many files in place can tune the thresholds. Alert-only incidents auto-expire after 7 days, blocks require a human to clear.
- **Storage blast-radius containment.** Every **paid** org gets **dedicated storage** by default: its own bucket, reached through a credential scoped to that bucket alone, so a credential minted for your org cannot reach another org's bucket. See [Dedicated Storage Isolation](#).
- **Data residency (jurisdiction pinning).** An org's metadata and dedicated block storage can be pinned to a real, infrastructure-enforced jurisdiction — EU or US — chosen once at signup and permanent thereafter; free and self-serve on every paid plan, not an Enterprise upsell. An individual project can also pin its **own** region at creation, independent of the org's: it gets its own metadata namespace and its own region bucket, so its content is never served from the org's storage. FedRAMP / FedRAMP-High storage regions are available on Enterprise, provisioned through sales (Swarmfile is not FedRAMP-authorized). This is a platform-level restriction, not an application-level convention: jurisdiction-pinned storage lives in a namespace invisible to requests that don't carry the matching jurisdiction. See [Data Residency](#).

For the strictest buyers: self-hosted control plane

Where "the vendor's infrastructure touches our data at all" is disqualifying (some government, defense, and critical-infrastructure work), the E2E tier may still not be enough because metadata is server-visible. For those cases a **self-hosted control plane** is the Enterprise option: the same code running on infrastructure you control, on a self-hostable runtime compatible with the platform the hosted service is built on. It is available on Enterprise today; a fully air-gapped deployment is on the roadmap — [contact us](#) to scope what it would involve for your environment. It is the top of the spectrum; the E2E tier is the middle; managed is the default.

Certifications and controls on the roadmap

- **SOC 2** is on our roadmap (Type I first, Type II thereafter); we have not yet begun a formal audit. In the meantime we answer security questionnaires directly and walk teams through this architecture on request.
- **SIEM forwarding** — a managed Splunk/Datadog/S3 stream of audit, access, and quarantine events — is planned. Today, a Pro-plan owner can export the [activity feed](#) to CSV (up to 100k rows per export), and [outbound webhooks](#) can push events to an endpoint you control on a periodic schedule, not in real time.
- **Legal hold & retention immutability** — a per-entry legal-hold flag overriding retention/trash purge — is planned.
- **DLP & egress controls** — watermarking, share-link domain allowlists, and export-policy gating — are planned.

If any of these is a hard requirement for you today, [contact us](#) to talk through sequencing. For the permission model that *is* live today, see [Permissions](#); for the E2E setup walkthrough, see [End-to-End Encryption Setup](#).

Trust & Compliance

A map for security reviews and procurement. Everything here links to the details; we claim no certification we don't hold.

Compliance posture today

Ask	Status
SOC 2	On the roadmap (Type I first, then Type II); no formal audit has begun . We answer security questionnaires directly and walk teams through the architecture on request.
FedRAMP	FedRAMP / FedRAMP-High storage regions are available on Enterprise, provisioned through sales rather than the self-serve picker. Swarmfile is not FedRAMP-authorized — talk to us before planning around it.
DPA	There's no published, pre-signed DPA today. Start the conversation via contact us and we'll tell you exactly what we can sign and work through your paper.
Subprocessor list	Below, plus the service-provider disclosure in the Privacy Policy .
Security questionnaire / security package	Answered directly, alongside the Security Architecture whitepaper.
EU / US data residency	Available today, free on every paid plan, chosen at org creation (or pinned per project at creation) — see Data Residency .
Bring-your-own storage	Available on Enterprise: your own S3-compatible bucket becomes the org's block storage — see Bring Your Own Storage .
Self-hosted control plane	Available on Enterprise: the same control-plane code on infrastructure you operate, on a self-hostable runtime compatible with the platform it's built on — see Deployment Topologies . A fully air-gapped deployment is on the roadmap.
SAML SSO	Available on Enterprise via an adapter; OIDC SSO is self-serve on Pro and above — see Identity .

Subprocessors

These are the third parties that process data for the hosted service:

Subprocessor	Role	What it can access
Cloudflare (Workers, R2, Pages, KV/D1)	Hosting for the hub, identity service, dashboard, and default block storage; optional Turnstile bot check at signup/contact; optional Workers Containers for video previews	Encrypted blocks for private projects; unencrypted for Free-plan/public projects by design; request metadata
Postmark	Notification and account email delivery	Recipient addresses and the notification content (event metadata, never file contents)
Stripe	Billing and payments	Account and payment details; never file content
iroh / NO relay network	NAT traversal for peer-to-peer connections when a direct path fails	Connection metadata; relayed traffic is end-to-end encrypted between peers
Google	Google Analytics on the public marketing/docs pages only (not loaded on public project pages, Explore, the signed-in dashboard, or app routes); Google Fonts site-wide	Page-view metadata (URLs, referrer, approximate location) from analytics; font requests carry the page URL as referrer. No project or file content
GitHub	Distribution of the git-LFS agent installer and release artifacts	Download metadata only

With [Bring Your Own Storage](#), block storage moves out of our hosted storage and into your own cloud account; with [Dedicated Storage](#), it moves into a bucket exclusively yours on our infrastructure. [Cloud-only mode](#) removes the peer-to-peer path entirely.

Data lifecycle commitments

- What's stored:** file content (encrypted according to the project's tier), block metadata, project history, and the activity/audit events described in [Telemetry & Data Collection](#).
- Retention:** history, trash, and audit rows default to **90 days** and are configurable 1–3650 days per org. Per-user preferences cover history and trash; the activity/ACL/audit rows follow the org window. Deletion audit rows for project and org purges are the exception: they are durable and outlive the purge, so a deletion stays auditable after the data is gone. See [Operations](#).
- Deletion:** canceling or an unresolved payment failure starts a **28-day grace period** with access blocked and reminder emails, after which the org's data and the owner's account (if it has no other org) are permanently deleted. The org owner can also delete an organization immediately from the dashboard. See [Data Portability & Offboarding](#).
- Export:** there is no lock-in format — the mount is a drive, so teams copy files out with normal tools; the activity feed exports to CSV; [Branch Mirror](#) continuously exports real files to a bucket you control. There is no indefinite post-cancellation read state, so export before you cancel.

Encryption you can point at

- At rest:** AES-256-GCM for private projects on paid plans, with per-project keys; Free-plan and public projects are plaintext by design, and some git-LFS uploads (direct stock-client uploads, and 64 MiB+ uploads through the built-in agent) are stored without application-layer encryption. The full tier comparison and exceptions are in [Security Architecture](#).
- End-to-end (opt-in, Pro+):** the key never reaches our servers; recovery is the customer's responsibility via [device transfer and recovery keys](#).
- In transit:** TLS to the hub and storage; QUIC with end-to-end encryption between peers.

- **Public release exposure:** a publisher's committed `.swarmfile/publicignore.yml` keeps matched paths and platform-derived data (currently CI-run logs) out of the anonymous release view, enforced as each release is published. Members see everything; share links are not covered. See [Security Architecture](#).

Security contact

- **Vulnerability reports:** the [security page](#) directs you to the contact form — put "Security vulnerability" in your message so it reaches the right people. Please include reproduction details and give us reasonable time before public disclosure.
- **Everything else (DPA, questionnaires, incident questions):** [contact us](#) and mark the request as a security review.

What is not available today

Procurement checklists often include these; we don't claim them today:

- SOC 2 report, penetration-test summary, or a public trust portal.
- A managed SIEM/streaming export (CSV export and webhooks exist — [Operations](#)).
- Legal hold / retention immutability, DLP watermarking, and share-link domain allowlists (roadmap).
- A fully air-gapped deployment (the self-hosted control plane exists on Enterprise; air-gapped operation is roadmap).

If one of these is a hard requirement, tell us during the review and we'll talk sequencing honestly rather than work around it.

Where to start a security review

1. [Security Architecture](#) — tiers, threat model, what each party can see.
2. [Security](#) — operational controls (encryption, integrity, quarantine, cloud-only mode).
3. [Network Requirements](#) — every host, port, and direction.
4. [Telemetry & Data Collection](#) — what is measured and what never leaves unencrypted.
5. [Data Portability & Offboarding](#) — exit path and deletion timing.

End-to-End Encryption Setup

Private projects on Starter and above are encrypted by default with a hub-held key ([Security](#) covers the model — Free-plan projects and public projects are stored unencrypted). End-to-end (E2E) encryption is a stronger, opt-in tier: the project key is wrapped to each member's own device keypair and never reaches Swarmfile's servers in plaintext, so the service cannot read your file contents at all.

Plan: Pro and above. E2E is enforced server-side — creating an E2E project on a plan without it is refused with a `402` (upgrade required), so pick the tier on a Pro or Enterprise org.

This page is the operational how-to — how to turn it on, set up recovery, and what to know before you do. For the cryptographic model and what E2E does and doesn't protect, read the [E2E section of Security](#) first; the short version is that it's **content-only** (filenames, folder structure, and sizes stay server-visible so search, ACLs, browse, and quota keep working).

Before you turn it on

Three things are worth internalizing before you create your first E2E project, because two of them are irreversible:

- **The tier is immutable per project.** A project is either managed-key or E2E, chosen at creation and never changed. You can't convert an existing managed project to E2E (the hub has already seen its key) or an E2E project back. To move existing data under E2E, create a new E2E project and migrate into it.
- **git-LFS is unsupported on E2E projects.** An LFS request returns a per-object `422` with a clear message: a stock `git-lfs` client can't hold a per-user key and the hub structurally can't read E2E content, so there's no one to encrypt or decrypt for. Use a managed or unencrypted project for LFS, or keep the asset on the mounted drive.
- **Server-side previews are disabled on E2E content.** Thumbnails, video posters, scrubbable proxies, and point-cloud previews are all generated server-side, which an E2E project's server can't do because it can't decrypt. E2E files show a generic icon in the dashboard. This is the single biggest day-to-day cost — see [File Previews](#).
- **"E2E with recovery" is not zero-knowledge-absolute.** If you set up an org recovery key (strongly recommended, below), a quorum of your admins *can* reconstruct a project key. That's a deliberate, enterprise-grade tradeoff — the alternative is that a lost password means permanently lost data. Choose consciously and tell your security reviewer plainly: recovery-enabled E2E means "Swarmfile can't read it, but your own admin quorum can."

We'd suggest piloting on a single non-critical project first: the tier is immutable per project, so a hasty rollout is expensive to walk back.

Step 1 — Set up the org recovery key *first*

Do this before creating any E2E project. When you create an E2E project, its key is wrapped to your org recovery key *if one already exists at that moment*. A project created before the recovery key exists has no recovery wrap — losing every member's device would then mean losing that project's data with no way back. Order matters.

As owner, open the **Recovery Key** tab in the dashboard (owner-only). Creating a recovery key:

- Generates an org recovery keypair and splits the private half with **Shamir secret sharing** into N shares with a t -of- N threshold (the default is 2-of-3). Any t shares can reconstruct it; any $t-1$ reveal nothing.
- Returns the share strings **exactly once**, on creation. They are never retrievable again. The view is print-friendly on purpose.

Distribute the shares to **different people and different physical locations** — a safe, a vault, an HSM, separate offices. The whole point of the threshold is that no single lost or compromised share exposes anything, and no single person can unilaterally decrypt. Store them the way you'd store the master keys to anything else that matters.

Rotating (creating a new key) or revoking is done from the same tab. Note that rotation applies to projects created *afterward* — it does not retroactively re-wrap projects created under a previous recovery key. **Keep the previous shares:** recovery accepts shares from any recovery key the org has created, active or previous, so those old shares are the only way to recover a project that was bootstrapped before the rotation. The dashboard warns about this at the rotate button.

Step 2 — Create an E2E project

Create a project as usual (**Projects** → **New**), and choose the **E2E** encryption tier instead of the default managed tier. The creation dialog states the tradeoff at the point of choice — previews disabled, recovery via the org key only, immutable after creation. E2E projects carry an **E2E** badge in the project list so the tier is never ambiguous.

Step 3 — Members enroll automatically

There's no manual key ceremony for ordinary members. When a member signs in on a machine, the engine enrolls that device: it registers the device's **public** key with the hub, while the private key is generated on the device and never leaves it. No admin action, no per-device setup screen.

Granting a member access to an E2E project happens **client-side**: an existing member's client wraps the project key to the new member's public key. The consequence worth knowing is timing — a newly added member sees an E2E project's contents only once *another* member's client has been online to complete that wrap. Add someone while the rest of the team is offline and their access is pending until someone with a key comes online. A running engine watches for the grant on its own and mounts the drive as soon as it lands — there is nothing to restart — and if that member switches to the project before the wrap is done, the switch waits for it rather than failing. If no grant was ever queued for them (for example a member added through a group), the engine asks the hub to open one the first time it tries to mount, so that wait heals without any manual step. An always-on **granter** closes this gap so grants don't wait on a human being online — enroll it yourself from the **Recovery Key** tab's "Granter" section (it needs a recovery key set up first, and reuses the same recovery shares): paste your threshold of shares once, and from then on new members and new devices get their keys automatically, on a five-minute cadence at worst.

If you'd rather not run a granter, or a grant is stuck waiting on one, any member whose own device already has the project's key can clear the queue by hand: open the project's **Permissions** tab in the web dashboard — a banner lists every pending device grant (who it's for) with a **Grant access now** button that wraps the key to each waiting device on the spot, and a **Dismiss** button per grant if it was created in error.

Multiple devices for one person

A member who works on more than one machine transfers their key between devices directly — a one-time, ten-minute code links a new device to an already-enrolled one and moves the wrapped project keys across, without the private key ever touching the hub. The code is bound to your account: a device signed in as a different account can neither link it nor fetch the moved keys. From the CLI: on the device that already has the project, run `swarmfile device-transfer start`; on the new device, run `swarmfile device-transfer link <code>` with the code it printed (see the [CLI reference](#) for the full details). The Desktop App's **Settings → Add another device** panel does the same thing without a terminal — generate a code on one device, type it into the other, and both sides poll and finish automatically once linked. Enrollment on the new device is otherwise the same automatic flow as the first.

Recovering a project key

If a member is locked out and their access can't be re-granted the ordinary way, an owner recovers the project key from the **Recovery Key** tab's recover flow: gather a quorum of shares from wherever you distributed them (for a project created before a recovery-key rotation, that is the *previous* key's share set — its threshold may differ from the active key's), enter them together with the project, and the dashboard reconstructs that project's plaintext key. This is a break-glass procedure — it requires the physical cooperation of your threshold of share-holders by design, and it's the reason the shares are stored apart in the first place.

Sharing E2E content externally

A [share link](#) on an E2E project still works for an external reviewer with no account: the key needed to decrypt travels in the link's URL fragment, which browsers never send to the server, so the hub serves only ciphertext and the recipient's browser decrypts locally. Through the link the hub limits raw-block access to the shared file's own blocks, so a recipient fetches only the file you shared — not the rest of the project. That limit is hub access control, not cryptographic separation: the link carries the *project* key, so treat sharing as a real trust decision and use expiry/revocation; revocation and per-email blocks stop block retrieval on the recipient's very next request. An access-count cap set on an E2E share through the API is enforced per distinct verified email — each new verified visitor consumes one access. (The deeper key-custody detail is in [Security Architecture](#).) Image and PDF previews work this way; video proxies don't (the server can't transcode what it can't decrypt), consistent with the in-app preview limits above.

When a member is removed

Removing someone from the org cuts off their access to E2E content within a few seconds: the hub stops serving them blocks and their wrapped project keys, revokes their key wraps, and cancels any key grants still pending for them. The revocation reaches every project in the org, including ones in the trash; a project that doesn't answer is retried until it does, and **Settings → Access revocations** shows any that are still pending. Their desktop drive stops serving files, clears its local cache of the project, and discards the project key it held. Changes they saved that hadn't finished uploading stay on their computer, still encrypted, and upload if their access is restored. A computer that's offline at the time does this when it next reaches the hub — the org's [offline access window](#), on by default at 7 days, bounds how long that can take.

What removal **can't** do is recall key material that's already on their device. A member who had access held a project key, and the hub can't reach into a machine to erase anything copied from it. If a departing member's copy of the key is a real concern, **rotate the project key** (below) so future writes stop using it.

Rotating an E2E project key

A project owner or admin can rotate an E2E project's key from the Desktop App's **Settings → My Devices** panel — the project's key-access card — and from the CLI with `swarmfile project rotate-key` (run it with the project open; it reads the active project off the mount). Rotation mints a fresh key generation and re-wraps it to every current device **and the org recovery key**, all client-side; the hub only ever stores the wrapped envelopes and never sees the key.

Know what it does and doesn't protect:

- **New content is protected.** Every write after the rotation is sealed under the new generation.
- **Existing content is not re-encrypted.** Blocks written before the rotation stay under the old generation, so a device that already unwrapped the old key can still decrypt content it already has or can still fetch from the swarm. Treat rotation as "stop using the old key going forward", not "make the old key useless everywhere".
- **Rotation is client-driven.** The hub never holds the key, so one of your devices performs it; if an org recovery key is configured, the new generation is wrapped to it too, so recovery keeps working. Only owners and admins can rotate.
- **Other devices keep reading.** A device that doesn't yet hold the new generation fetches it on its next check and keeps its old-generation key so existing files still open.

Seeing and revoking a project's devices

Owners and admins can see which devices hold an E2E project's key. Open the project's **Permissions** tab in the web dashboard and look under **Devices with access**. Each device is listed with its name, the person it belongs to, its platform, app version and when it was last seen. It also shows whether it still has access and, if not, why: removed from the organization, revoked by an admin, revoked by its owner, or waiting for a new grant.

Revoke cuts one device off from that one project. The device stops opening the project's files, and the copy cached on it is cleared when it next connects. It doesn't get access back automatically, and the person's other devices aren't affected. You can also revoke a device that has lost access but would be given it again when it next asks (removed from the organization, or waiting for a new grant), so a lost device in those states stays locked out. As with any revocation, anything already copied off the device can't be recalled. To cut a device off from every project at once, its owner can revoke it under **Settings → My Devices**.

What lives where

For a regular member, enrollment and key-granting are **automatic**: they happen in the engine without user action, and there's no Desktop App UI for them (nor anything to configure there). The owner-facing admin — the recovery key, its shares, and the recover flow — lives in the web dashboard only.

A member's own device keys, though, are self-service from the web dashboard, the CLI, and the Desktop App — distinct from the owner-only recovery-key/quorum machinery above, which stays web-only. In the dashboard, **Settings → My Devices** lists every device you've enrolled and lets you revoke a lost one yourself, no admin needed. From the CLI, `swarmfile e2e-key status` shows this device's own enrolled key (id, public key) straight from its local cache, no hub round trip; `swarmfile e2e-key list` audits every device enrolled on your account, and `swarmfile e2e-key revoke <id>` kills a lost one's access the same way the dashboard does.

Revoking a device takes effect on every E2E project in the org. If the device is online, it's cut off within seconds: its drive stops serving the org's E2E projects, and it discards the project keys it held along with its local cache of their files. If it's offline, the same happens as soon as it reconnects. It's also never granted a project key again, even from a request made on one of your other devices. Revoking a device key doesn't sign that device out, though. Revoking the key can't recall a project key that was copied off the device before you revoked it.

The Desktop App's **Settings → Add another device** panel walks through the same transfer `swarmfile device-transfer` does below, without touching a terminal.

Headless/CI identities work differently again: `swarmfile api-key create` mints a project-scoped API key, and for an E2E-tier project that command *is* a key-management operation — it generates real E2E key material (a device keypair, project-key wrap) for that key to use, printed once alongside the key itself. `swarmfile api-key revoke` reverses it, tearing down that identity's key material along with the key. There's still no Desktop App surface for either of these CLI paths.

The bottom line for a security review

- Contents are unreadable to Swarmfile; metadata (names, structure, sizes) is not.
- Recovery is your own admin quorum, not the vendor — and only if you set the recovery key up before creating projects.
- Previews and server-side proxies are off for E2E content, by necessity.
- Set-up is irreversible per project; pilot before you commit a flagship project to it.
- Removing a member revokes their key wraps and stops their access, but can't recall a project key already on their device.
- An E2E project key can be **rotated** (client-side, owner/admin) so future writes stop using an old key — but that does **not** re-encrypt existing content, so a device holding the old key can still read what was written before the rotation.

If E2E is a hard procurement requirement, raise it early and we'll walk your reviewer through the envelope format and the recovery threat model directly.

Dedicated Storage Isolation

Dedicated storage is the default for every new paid organization: your org gets its own bucket, reached through a credential scoped to that one bucket. The credential your org's traffic uses cannot touch any other org's data, even if that credential is compromised or misused. Free-plan orgs, and paid orgs that explicitly opted out at creation time, use shared storage: one shared bucket in which every org's objects are scoped under its own prefix.

This page is the operational how-to. For the underlying model — what "blast radius" means here and why the boundary is the org rather than the project — see [Security Architecture](#).

What this does and doesn't change

- **Bucket isolation, not encryption.** Dedicated storage is about *where the bytes live*, not what protects them once there — it gives your org a storage credential scoped to its own bucket alone, so it can't reach another org's data. Encryption at rest is a separate question with its own tiers: private projects on paid plans are encrypted; Free-plan and public projects are stored unencrypted (see [Security](#)).
- **Still our infrastructure, not yours.** This is **not** bring-your-own-storage — your data stays in a bucket we create and operate on our own cloud infrastructure, just a bucket that's exclusively yours instead of shared. If your compliance requirement is specifically "our data must live in infrastructure we control," this isn't the right tier: for that, [Bring-your-own-storage](#) (Enterprise) makes a bucket on *your* own account the org's primary block storage, and a self-hosted Enterprise control plane (see [Deployment Topologies](#)) puts the control plane and storage both on infrastructure you control. Talk to us about which fits.
- **The org boundary, not the project boundary.** One dedicated bucket covers every project in your org. Projects within it are still separated the same way they always were (a project-scoped prefix inside the bucket) — that separation was already correct and doesn't need its own bucket to be safe.

Requesting it

New paid orgs select this automatically — there's nothing to request. Org creation records the storage choice but does not provision a bucket. The dedicated bucket is created only when the user deliberately creates the org's first project. API callers can also explicitly opt out at creation time with `POST /orgs` and `blockIsolation: "shared"`.

If your org was created with the shared opt-out and you've changed your mind, an owner can still turn dedicated storage on from **Settings → Storage** — a one-click, one-way switch (see below).

Starter plan or above only. Free-plan orgs always use shared storage and aren't offered the switch at all; upgrading to Starter or Pro is the path onto dedicated storage, and the dedicated bucket is then created with the org's first project (or by the switch itself, if the org still has no files). The same rule is enforced server-side, not just hidden in the UI.

This only works before the org has any files. There's no migration path yet to move existing files into a new bucket, so switching is rejected outright once an org has uploaded anything. For an org that opted out at creation and wants to reconsider, that means: do it before adding content — not later as an afterthought.

The same page also carries your bucket's public-access switches when your org is on [Bring-your-own-storage](#) — anonymous streaming and gated release downloads served from a customer-supplied bucket. Shared and Swarmfile-managed dedicated buckets always stream and serve downloads, so those switches only appear in BYOS mode.

What happens during provisioning

Provisioning isn't instant — creating the bucket and minting a credential scoped to it both involve real calls to our storage platform, and a freshly minted credential needs a few seconds to become valid before we'll rely on it. This begins when the first project is explicitly created, or on-demand when an existing eligible org turns the switch on from Settings → Storage. Until it finishes, the org's block storage briefly denies reads and writes rather than falling back to the shared bucket — a silent fallback would defeat the isolation guarantee, so a short, visible pause is the deliberate tradeoff. If provisioning hits a transient failure (an upstream API hiccup, for instance), it retries automatically with backoff, so nothing needs to be re-requested by hand.

Known limitations

- **Empty orgs only.** No migration path exists to move already-uploaded files into a new dedicated bucket — the switch is rejected if the org has any files. This isn't a request-and-wait limitation; it's enforced at the moment you ask. A Free-plan org that upgrades to a paid plan *after* uploading anything is in the same position: the upgrade gets dedicated storage for future work only when the org was still empty, and an org that already has files stays on shared storage.
- **One-way.** There's no supported path to move an org that's already on dedicated storage back to the shared tier — the Storage page's switch only goes one direction, on purpose. Make sure this is what you want before turning it on.
- **No self-service credential rotation.** If you have a specific reason to believe your org's storage credential needs rotating, contact us directly — there is no self-service flow.
- **Owner-only.** Admins and members don't see the Storage tab; only the org owner can turn it on — and a Free-plan org's owner can't turn it on at all (see above).

For API consumers

`GET /orgs/:orgId/block-bucket` reports `hasContent` — whether the org already has files, the fact that decides eligibility for the one-way switch. It is only meaningful for a `shared` org and is computed with a cross-project lookup, so the hub skips that lookup and reports `false` for `dedicated`, `byos`, `provisioning`, and `failed` orgs; don't read `false` as "no files" in those modes. `mode` is the field to branch on everywhere.

Bring Your Own Storage

Bring-your-own-storage (BYOS) makes a customer-supplied S3-compatible bucket the **primary** block storage for an organization. Instead of your data's blocks living in a bucket on our storage platform, they live in a bucket on *your* account, on *your* provider (AWS S3, Wasabi, Backblaze B2, MinIO, and other S3-compatible targets) — with our hub still running the control plane, metadata store, and coordination as usual.

BYOS is an **Enterprise-only** option (the `byos_storage` entitlement).

How it's different from the alternatives

- **Not Dedicated Storage.** [Dedicated Storage](#) gives your org its own bucket *inside our own infrastructure* — physically isolated, but still a bucket we create and operate on our cloud account. BYOS moves the bucket out to infrastructure *you* own and control.
- **Not a branch mirror.** A [branch mirror](#) is a read-only export *copy* of one branch into your bucket, alongside the authoritative data that still lives on our platform. BYOS is the opposite relationship: your bucket *is* where the live blocks are read from and written to — there's no second copy on our side.
- **Not full self-hosting.** BYOS relocates only the block storage. The hub, metadata store, and coordination still run on our infrastructure. If your requirement is that *everything*, control plane included, runs on infrastructure you operate, that's the self-hosted Enterprise control plane ([talk to us](#); fully air-gapped operation is on the roadmap), not BYOS.

Think of it as a spectrum of "where do the bytes live": shared bucket (Free-plan and shared-storage orgs) → your dedicated bucket on our infra (Dedicated Storage, the default for new paid orgs) → **your bucket on your infra (BYOS)** → your entire deployment on your infra (self-hosted).

Requesting and provisioning it

BYOS is provisioned by the **org owner** through Enterprise onboarding. Because it makes an external bucket the org's primary storage, it's a deliberate, owner-gated action — not a self-serve toggle a member could flip.

You supply a bucket that **already exists** plus a credential that can read and write it; Swarmfile validates the credential with a real write-read-delete round-trip against your bucket before switching anything over, so a bad endpoint, region, bucket, or key is caught immediately rather than surfacing on the org's first real read or write.

Configuration fields

Field	Notes
Endpoint URL	Your provider's S3-compatible endpoint.
Region	The bucket's region.
Bucket name	The target bucket. It must already exist.
Key prefix	Optional sub-path inside the bucket to confine Swarmfile's objects to (e.g. <code>swarmfile/</code>).
Access key ID	The S3 access key ID for a credential with read/write access to the bucket.
Secret access key	Stored encrypted at rest (AES-256) under a dedicated key, and never returned.

The endpoint must be an `https://` URL on a real remote host. A cleartext `http://` endpoint would put your secret key on the wire, so it's refused unless your deployment operator has explicitly opted in for a trusted on-prem target; loopback, link-local, and cloud-metadata addresses are always refused. A self-hosted operator can additionally pin which hosts are allowed.

Content is encrypted at rest exactly as it is on our own storage for **private, managed** projects (see [Security](#)) — the objects that land in your bucket are the same encrypted blocks, just hosted by you. **Public projects are unencrypted by design** (they have to be, to be served anonymously), so a public project's blocks land in your bucket as plaintext.

Public projects: streaming and downloads from your bucket

If your org publishes a public project, visitors can play its video and audio and view large images straight away, without downloading the whole file first. For a BYOS org that streaming is served **from your bucket**, so the requests and the data transferred out are billed to **your** storage provider account, not ours.

Two independent owner switches control this, both on **Settings, Storage**, and both take effect on the next request:

- **Anonymous streaming (Allow anonymous streaming from my bucket).** Because anyone on the internet can trigger those reads, this is **off by default**.
 - **While it's off:** public visitors can't stream media from your bucket. On the public project page the audio, video, or image falls back to the sign-in gated download, which needs a Swarmfile account. Thumbnails, small previews, and readmes are not affected, and neither is anything for signed-in members of your org.
 - **While it's on:** anyone who can see one of your public projects can stream its media directly from your bucket. Expect your provider to bill for the requests and egress. Swarmfile applies per-IP rate limits and byte budgets to public streaming, but it cannot cap what your provider ultimately charges you.
- **Gated release downloads (Allow anonymous release downloads from my bucket).** Release files and a release's combined archive, already behind a sign-in, a solved proof-of-work challenge and rate limits. For a **new** BYOS org this is **off by default**, and gated release downloads are refused until you turn it on — visitors see the same "not found" a private release shows. An org with existing public releases defaults to on so those releases keep working.

Shared and Swarmfile-managed dedicated buckets stream and serve downloads by default; the switches only appear in BYOS mode.

Known limitations

- **New/empty orgs only.** BYOS is set at provisioning time, before the org has uploaded any content. There is **no live migration** of an existing org's blocks into a BYOS bucket in this version — switching an org that already holds data is rejected. Decide on BYOS before adding content, not after.

- **No fallback if your bucket is unavailable.** This is the real operational tradeoff to understand. Your bucket is the *only* home for the blocks — there's no shadow copy on our side to fall back to (a silent fallback would defeat the entire point of BYOS). If your bucket has an outage, or its credential is rotated/revoked out from under Swarmfile, reads and writes for the org **fail synchronously** until you restore access. Storage availability and credential lifecycle become your responsibility.
- **Credential rotation is re-provisioning.** When you rotate the access key, re-run provisioning with the new credential for the *same* bucket. (Re-pointing a non-empty org at a *different* bucket is refused for the same reason a live migration is — it would orphan the existing content.)
- **Set the incomplete-multipart lifecycle rule on your bucket.** A large git-LFS push (≥ 64 MiB, via the desktop app's transfer agent) is one S3-compatible multipart upload that goes **straight to your bucket**; a client killed mid-push can't abort its own upload, and your provider bills the uploaded parts until they're aborted or expired. Add an `AbortIncompleteMultipartUpload` rule — 7 days is the value we use. Swarmfile applies this automatically to buckets it creates; a BYOS bucket is yours to configure.
- **Data residency is yours to enforce.** Swarmfile can't verify the physical region of an arbitrary customer endpoint, so BYOS is refused for an org with a [pinned jurisdiction](#) — and for an org that has any project pinned to its own region. If you have a residency requirement, choosing a bucket in the right region — and keeping it there — is on you.
- **Enterprise-only, owner-gated.** BYOS requires the Enterprise plan, and only the org owner can provision it.

Where to go next

- [Dedicated Storage Isolation](#) — physical isolation inside our own infrastructure.
- [Branch Mirror](#) — a read-only export copy into a bucket you control.
- [Deployment Topologies](#) — the fully self-hosted Enterprise option.

Data Residency

Data residency lets you pin an organization's metadata and block storage to a real, infrastructure-enforced jurisdiction — **EU** or **US** — chosen once at signup. It's free and self-serve on every paid plan, Starter through Enterprise (Free-plan orgs use shared storage and can't pin a region); there's no price gate and no sales conversation needed to turn it on. An individual project can additionally pin its **own** region at creation, independent of the org's — see [Per-project residency](#). FedRAMP and FedRAMP-High storage regions are available on Enterprise, provisioned by us through sales rather than the self-serve picker. Swarmfile itself is not FedRAMP-authorized — the regions are available, the authorization is not — so if you have that requirement, talk to us before you plan around it.

This page is the operational how-to. For the underlying mechanism — how jurisdiction pinning actually works at the infrastructure level — see [Security Architecture](#).

What this does and doesn't change

- **Platform-enforced, not a convention.** This isn't an application-level rule we promise to follow — it's a real, platform-level restriction on where your organization's metadata and block storage physically live. Jurisdiction-pinned storage exists in a namespace that's invisible to any request that doesn't carry the matching jurisdiction; there is no shared code path that could accidentally route around it.
- **Covers our system of record, not your own peer-to-peer transfers.** Your own team's machines transferring files to each other over the peer-to-peer engine are already authorized devices on your own network — that's not something jurisdiction pinning governs, the same way a cloud region guarantee doesn't follow your own laptop once you've downloaded a file. If you have your own downstream residency commitment to your own clients, you can additionally turn on cloud-only mode (Settings → Network) to disable peer-to-peer entirely for that org.
- **Global by default.** If you don't choose a jurisdiction at signup, your org runs on our standard global infrastructure — this is a purely additive choice, not a change to what already exists.
- **Immutable, on purpose.** Jurisdiction is fixed at creation by design. Both your metadata store and your block storage have their jurisdiction fixed at creation by the platform itself; moving one means migrating to entirely new infrastructure, not flipping a setting. Choose carefully at signup. The same applies to a per-project pin (below).

Choosing it

Pick **Data residency** on the signup form (or the organization-creation screen if you're joining as an existing user with no organization yet) — the choice is right next to the organization name field. Leaving it on **Global (default)** is the same as not having this feature at all; picking **European Union** or **United States** pins the org immediately upon creation, and a confirm dialog gates the choice specifically because it can't be undone.

FedRAMP and FedRAMP-High aren't in the self-serve picker, but they are available on Enterprise: we provision them through sales. If you have a FedRAMP requirement, talk to us directly.

Per-project residency

A paid org can pin an **individual project** to its own EU/US region when the project is created — independent of the org's own choice, so a project on a Global org can be pinned, and a project on an EU org can be pinned to US. A pinned project gets its own metadata namespace and its own region bucket, and is never served from the org's storage.

- **Chosen at creation, immutable.** Like the org-wide choice, a project's region is fixed when the project is created; there is no way to move or un-pin it later. The create form states this before you commit.
- **Entitlement.** Per-project EU/US pinning is available on every paid plan; FedRAMP regions are Enterprise, provisioned through sales (not the self-serve picker), and Swarmfile itself is not FedRAMP-authorized.
- **Provisioning.** The region's bucket is provisioned on first use. Until it's ready the project's reads and writes fail closed with a `storage_provisioning` refusal — it never silently falls back to the org's bucket — and the creator is notified once it's ready.
- **Rules out customer-managed storage.** A pinned project can't use a [branch mirror](#), and an org with any pinned project can't switch to [BYOS primary storage](#): Swarmfile can't verify the region of a customer-supplied endpoint.
- **Web, desktop app and CLI.** `swarmfile project create --residency eu` pins at creation, and `swarmfile project show` reports the project's region and storage status (see the [CLI reference](#)).

Known limitations today

- **No self-serve upgrade path.** If your org is already running on global infrastructure, there's no way to move it to a pinned jurisdiction after the fact. This is the same "provisioning is a one-time, creation-time decision" constraint as [Dedicated Storage](#), for the same underlying reason — neither your metadata store nor your block storage can be relocated once created.
- **FedRAMP is Enterprise, not self-serve.** There's no self-serve picker for it — we provision FedRAMP/FedRAMP-High regions through sales — and Swarmfile is not FedRAMP-authorized; ask us before planning around it.

Billing & Plans

How billing works

Signing up for **Starter or Pro** runs through Stripe Checkout. Once you're on a plan, you can **switch between Starter and Pro right from the Billing tab** — "Upgrade to Pro" / "Switch to Starter" — without leaving the app. The switch retargets your existing subscription in place (it never creates a second one), takes effect immediately, keeps any trial in progress, and confirms the projected cost first. Canceling, updating a payment method, and viewing invoice history still happen in a self-serve Stripe Customer Portal, available to the org owner (which also offers the same plan switch as a fallback). The Billing tab shows a **"What happens if you cancel?"** summary next to the portal link — 24 hours, then the 28-day grace period with access blocked, then permanent deletion — so the consequences are visible before the cancellation button, not only after. See [Data Portability & Offboarding](#) for the full detail.

Free is different: it's a no-card tier you choose at signup, not something you switch into or out of later. There's no Stripe subscription behind it at all — nothing to charge, cancel, or retarget. Upgrading from Free runs the same org through Checkout for the first time, the same as any Starter/Pro signup; there's no separate downgrade path back into Free afterward.

One thing neither the app nor the portal lets you edit **once you're subscribed**: seat *quantity*. (At sign-up you can pre-pick a starting seat count — that's the number the first invoice uses — but from then on it's live.) Seat count is derived from org membership, not a number you type — **you change your seat count by adding or removing people in the Team tab**, and your billed seats follow automatically, reconciled daily. Editing a seat count by hand would just get overwritten by the next reconcile sweep; the Billing tab links you to Team instead.

Starter and Pro carry the same 7-day free trial. Your card is collected upfront at signup, and the first charge happens on day 7. **During the trial, seats are capped at 2 and add-ons aren't available** — growing past 2 people, or adding extra storage/collaborator capacity, means ending the trial and paying first (you can do that any time, see below). Trials are also one-time per org: once a subscription has gone active, canceling and resubscribing doesn't grant a new trial period.

The trial is optional, and you can leave it early:

- **Skip it.** At sign-up, tick **Skip the 7-day trial and start paying today**, or pick more than 2 seats (a trial can't cover them, so that starts paid automatically). From the Billing tab, Starter and Pro each offer **Subscribe now, skip the trial** next to the trial button. Either way you get a paid subscription straight away, with no trial and no 2-seat cap.
- **End it early.** While trialing, the Billing tab's **End trial and subscribe now** charges your card on file for about the first month's seats right away (the confirmation shows the amount) and lifts the 2-seat cap immediately; add-on packs unlock at the same moment. If the card is declined, nothing changes: you stay on the trial, and you can update the card in the Customer Portal and try again. If you canceled during the trial, resume the subscription in the Customer Portal first; the Billing tab points you there instead of offering to end a trial that's already set to stop. The Team tab's seat-limit notice and the desktop app's trial notice both link there, and the Billing history records the change as **Trial ended early**.

Plans

Plan	Price	Storage	Seats	External collaborators	Headless API keys
Free (community in the API)	Free, no card	Grows with account age, up to 1 GiB total	1 seat only	Free on public projects	1
Starter	\$5/seat/mo	100 GiB/seat	1+ seats, no ceiling	1 per seat (private projects); public projects free	2 per seat (min 5)
Pro	\$25/seat/mo	500 GiB/seat	1+ seats, no ceiling	5 per seat (private projects); public projects free	5 per seat (min 10)

The seat figures above describe an active, paid subscription — the 7-day trial itself is capped at 2 seats regardless of plan (see above). The Free plan has no trial (there's nothing to trial) and no add-ons; its storage cap is a single all-in allowance that starts small on a brand-new org and rises to 1 GiB over its first few months, rather than the per-seat allowance Starter and Pro use.

The Free plan can only create **public** projects, and every project it creates is stored **unencrypted** at rest — private projects and managed encryption at rest both require Starter or above; the end-to-end tier requires Pro or above. The same is true of any **public** project on a paid plan: public projects must be unencrypted so they can be served anonymously. These are the two things the Free plan restricts by *kind*, not just by size: a Free-plan org that wants either has to upgrade first, it can't simply run out of room the way storage or collaborators can. Encryption tier is also fixed at project creation — upgrading doesn't retroactively encrypt a project created while on Free; only projects created after the upgrade get managed encryption. See [Security Architecture](#) for the full detail.

Each plan also caps how many **live projects** an organization can have: 10 on Free, 100 on Starter, 250 on Pro (Enterprise is contract-negotiated). Deleting a project frees its slot; the cap bounds per-org resources and never bills anything.

Everything stored at rest counts toward your storage allowance on the same basis — project files, and the objects held for a project used as a [git-LFS](#) server (billed on distinct content, so deduplication applies). Storage starts with the first project: an org with no projects has no billed storage and no dedicated bucket yet. Bandwidth is unlimited and free on every plan, for both cloud and P2P transfer. Every plan's included bandwidth allowance is unlimited: egress is never billed and no plan caps it, and egress bytes are still counted per org and shown on the Usage tab. The one ceiling is the platform-wide anonymous-read abuse backstop (see [Telemetry & Data Collection](#)), which can refuse anonymous public reads at extreme volume. Unlike storage and collaborators, egress has no billed meter at all, so "unlimited" is a policy choice on top of real counters, not an absence of counting. Self-hosted seed/NAS nodes require Pro or above; see [Self-Hosted Seed Nodes](#) for setup.

Need more than Pro's limits — unbounded external collaborators and headless API keys, a custom seat arrangement, or different billing terms? Enterprise plans are available; contact us.

Overage rates

Storage that exceeds your plan's included allowance is billed at \$4/100 GiB on Pro and \$5/100 GiB on Starter. External collaborators beyond your plan's included allowance are billed at \$1/month per collaborator, up to the 3x hard ceiling described below — **private-project grants only**. Collaborators on public projects are free and consume nothing: neither the included allowance, the overage meter, nor the ceiling. The Free plan bills no overage on anything — there's nothing to charge a

card that doesn't exist — so its storage allowance is a hard wall, not a soft one (see Enforcement below). Headless API keys never bill overage — they're a hard cap, not a metered dimension: Starter includes 2 per seat (minimum 5) and Pro 5 per seat (minimum 10), and committed key packs add 10 keys for \$20/month on either plan (see Enforcement below).

Prefer committed capacity to overage? **Storage packs** pre-buy it at a discount: **\$4/100 GiB on Starter and \$3/100 GiB on Pro** per month — a dollar below the matching overage rate. A pack adds 100 GiB to your included allowance for as long as it's on the subscription, and because the allowance is raised first, pack capacity is never also billed as overage (and never counts against the spend cap below). Manage them in the **Add-ons** section of the Billing tab; a plan caps pack counts per kind — 100 per kind on Starter; on Pro, 10,000 storage packs or 1,000 collaborator/key packs — because each committed pack also lifts the 3x hard ceiling. Collaborator packs work the same way: \$0.80/month per extra external collaborator, versus \$1 in overage. Changes are also rate-limited (about ten per organization per day), and an increase must clear its prorated charge immediately rather than waiting for the next invoice.

Hosted CI compute is a third metered dimension, separate from the allowances above: **Hosted CI** jobs bill per second of container time at our underlying provider's at-cost rate × 1.1, with a 60-second minimum per job and no included allowance on any plan — every cent counts against the spend cap from the first second. It is deliberately not plan-gated: every plan can enable it per project. A Free-plan org has no card by default, so it attaches one for compute only (a zero-price subscription, not a plan upgrade) before its first job, and then runs under a platform compute ceiling scaled by account trust.

Enforcement

Quota enforcement is always on. There's no setting to disable it and no feature flag gating it off — it's a structural part of the product, not an opt-in.

Plan storage counts the size of your files. Each project also has a separate, unbilled limit on its file information and history (names, versions, and so on), which only very large projects approach — see [Project storage limit](#).

Storage, seats, and external collaborators are enforced differently, on purpose:

- **Seats** have no ceiling and no minimum on either plan once you're subscribed, beyond the universal 1-seat floor every org has (you always have at least an owner). Add or remove people in the Team tab any time; billed seats simply follow. The one exception is the 7-day trial itself, which is capped at 2 seats until you commit to a plan (see [Plans](#) above). The Team tab reflects whichever cap applies — it replaces the invite form with the plan's limit and an upgrade link at the Free plan's single seat, and a **Subscribe now** link (to end the trial) at the trial's two, so the refusal is visible before anyone types an address.
- **Storage and private-project external collaborators** both bill metered overage before hard-blocking, rather than blocking right at the included amount: storage hard-blocks at 3x the included allowance, and private-project external collaborators hard-block at 3x the included per-seat allowance. (Public-project collaborators sit outside this meter entirely — free, no slot, no ceiling.)
- **Headless API keys** are a hard cap, not a soft one: minting is refused with a `402 (code: keys_cap)` the moment the org is at its allowance — Starter 2 per seat with a floor of 5, Pro 5 per seat with a floor of 10, Free 1 — and key packs (+10 keys for \$20/month) raise the cap. Keys are org-owned machine credentials for headless fleets (render nodes, CI, containerized agents, seed/NAS nodes); user workstations and their mounted agents need none, and personal access tokens (a person's own automation) are free. Revoking a key frees a slot immediately, and a downgrade never revokes keys you already have — it only blocks minting new ones past the new allowance. A key is walled off to exactly one project, and deleting a project revokes its keys with it. Owners see the standard soft-cap warning banner in Billing once the org reaches 80% of its key allowance, before minting ever refuses.
- **The Free plan's storage** is the one exception to the above: its included allowance and its hard block are the same number. There's no overage zone to sell, so there's nothing to leave room for — the cap you see is the cap that's enforced.

The reasoning: the pricing page sells metered overage as a real capability for both dimensions, so hard-blocking the moment you cross the included allowance would refuse to sell what's advertised. The backstop exists to catch runaway or abusive usage, not to cap normal overage billing. The Free plan isn't selling anything, so this reasoning doesn't apply to it.

Spend cap

Starter and Pro also carry a **spend cap**: a monthly dollar ceiling on metered charges (storage above your allowance, extra external collaborators, and hosted-CI compute — metered from the first second, with no included allowance). Seats and add-on packs aren't counted against it, and downloads are always free. By default the cap equals your plan's seat charge — your seat count times the per-seat price — and an owner or admin can change it under **Billing → Spend cap**, either permanently or for 24 hours, 7 days or 30 days before it reverts to the default. Leaving the field blank, or choosing **Reset to plan default**, clears your own figure. Resetting to a default that sits below metered charges already accrued this month pauses new uploads immediately, exactly like lowering the cap by hand — the form says so before you click. There's no upper limit on the figure: above the most storage and collaborator growth alone could bill (the 3x hard ceiling), metered compute can still push usage past the cap, and the Billing form says so as you enter it.

The cap is a hard ceiling on the overage line itself, not just a pause button: overage bills up to it and never past it. Because metering follows your *standing* usage for the period, the ceiling applies to the whole month — lowering it below overage you've already accrued drops that month's charged overage to the new figure and pauses further growth immediately, while raising it lets accrual continue. It never changes what's included: plan seats and committed packs still raise your allowance before any overage is counted.

When metered charges reach the cap, uploads that would add billable storage (and invites that would add a billed collaborator) pause, and new hosted-CI jobs are refused with a blocked `spend_cap_reached` run — running jobs finish. Reading, streaming, deleting, renaming and everything else keep working, and nothing is lost: the desktop app holds paused saves and uploads them once the cap is raised or storage is freed. A refusal is a `402` with `code: spend_cap_reached` carrying `capCents`, `usedCents`, `remainingCents` and the figure the request would have reached; the desktop app shows "Spend cap reached" and keeps retrying, so a raise resumes the backlog on its own. The Free plan and Enterprise have no spend cap on the storage/collaborator line — a Free-plan org's hosted CI runs under the platform's Community compute ceiling (after a card attach), and Enterprise compute is contract-priced.

A payment **dispute** (a chargeback on the subscription) is a different, deliberately sticky lockout rather than a cap: while it is open, the org's file reads and writes are refused with a `402 (code: grace_period_inaccessible)` until the dispute is resolved with support. Nothing is deleted, and saves that were paused upload as soon as it is cleared.

Per-key budgets. A project-scoped API key can also carry its own ceiling, set per key on the **API Keys** page (blank means no key budget). Enforcement uses the lower of the two — `min(org_spend_cap, key_budget)` — and a key-budget refusal affects only that key's writes: other keys, user workstations, and reads are untouched. A key refusal names the key in the `402 (keyLimitCents, keyId)`.

Notifications. Owners and admins get an in-app notification when the cap is approached (50%, 80%, 100%), on the first refusal of a billing period, when a raise or a cleanup lets writes resume, and whenever the cap or a key budget changes. Billing's "approaching plan limits" banner shows the same levels.

Plan-gated features

Three features are **Starter and above**: private projects, a dedicated per-org storage bucket (the one-way enable on **Settings → Storage**), and an [organization-enforced authentication policy](#) (require MFA and/or a verified email for members) — a Free-plan org can only create public projects, stored unencrypted (plaintext), and can't turn a policy requirement on; requesting a private project (managed encryption at rest comes with it) from the Free plan, or enabling a policy requirement, is rejected the same way any other plan-gated feature is. A number of others are **Pro and above**: SSO (OIDC), SCIM/LDAP sync, folder ACLs with Windows DACL projection, audit-log export, the [end-to-end-encryption tier](#), [self-hosted seed nodes](#), and [branch-mirror-to-S3](#). A few are **Enterprise-only**: [bring-your-own primary storage](#), SAML SSO (via an adapter — see [Identity](#)), and a [self-hosted control plane](#). (EU/US [data residency](#) is the exception — available at no extra cost on every paid plan, not an Enterprise upsell; a Free-plan org uses shared storage and can't pin a region.) These are enforced server-side with a real `402` rejection if you try to configure them on a plan that doesn't include them — not just hidden behind UI. Downgrading never disables something you've already configured; it only blocks configuring something new.

For the details on those features themselves, see [Identity](#) and [Permissions](#).

Usage visibility

Plan standing, current usage, and your org's resolved entitlement set are visible to any org member, not just the owner. See [Organizations](#), [Projects & Members](#) for where that lives. The **Usage** tab (owners) breaks usage down by dimension — storage, egress, collaborators, hosted-CI compute — and includes a **headless API keys** line showing how many of the plan's included keys the org is using; Billing carries the 80% warning banner before minting is refused.

Operations

This page covers the tools you'll reach for as the person operating an org day to day: the audit log, retention policy, and connectivity diagnostics.

Audit log

The Activity tab in the web dashboard is a unified activity feed, open to any org member, covering (among others):

- ACL changes
- History events
- Quarantine incidents — both applied and released
- Branch-protection changes — per-branch flag toggles and glob protection-rule edits
- Org-level audit events (including member role changes and data-plane policy changes)
- Notifications
- Unlock requests
- Changeset commits
- Branch creates, archives, and switches
- Tag creates
- Comments added, resolved, and mentions
- RFI / submittal / approval workflow events (created, assigned, due-soon, overdue, response posted, revisions requested, closed, voided)
- Webhook changes (created, edited, rotated, re-enabled, or removed)

The feed is ACL-filtered server-side: owners see everything, members see events on entries they can read, and quarantine incidents and org-policy events project to empty for non-owners. Everything lands in one feed instead of being scattered across separate logs. It refreshes automatically (within about a minute while the tab is open, paused in background tabs) so new events show up without reloading the page. You can export it to CSV, capped at 100k rows per export. The export endpoint is owner-only *and* gated on the `audit_log` entitlement (Pro plan and above): a non-owner gets a 403, and an owner on a plan without the entitlement is told to upgrade. The live in-dashboard feed and the underlying event capture stay ungated on every plan — only the bulk CSV export is entitlement-gated.

Today the underlying events are stored internally, with CSV as the only bulk-export path — though an [outbound webhook](#) does let you forward a live stream of the same underlying events to your own endpoint (delivered on a periodic schedule, not in real time), which covers some of the same ground for a system willing to receive and store them itself. See [Security](#) for the full list of what's still on the roadmap.

What the CSV export contains

Each export starts with a few `#` provenance rows (format version, generation time, requested window, row cap) that spreadsheet apps skip, then one row per event with these columns: `occurred_at_iso`, `kind`, `actor_user_id`, `actor_display_name`, `actor_email`, `project_id`, `project_name`, `entry_id`, `entry_name`, `summary`, `detail`. The rows are a merged, deliberately un-sorted-across-sources view of every source in the requested window; sort by `occurred_at_iso` after import if you need a single timeline, and check the `X-Activity-Export-Truncated` response header in the web export. The CLI prints the CSV without surfacing that header, so a CLI export is silently capped at 100,000 rows — narrow `--since / --until` if you need a complete window.

Audit-event retention. Activity, ACL, and audit rows are pruned on the org's retention window — the same 1–3650 day setting that governs history and trash, defaulting to 90 days when nothing is set. Per-user retention preferences apply to history and trash only; these activity/ACL/audit rows follow the org window. Plan for exports or webhook ingestion if your compliance position needs records kept longer than the window in force, and confirm the window on **Settings → Organization** before relying on it.

Ransomware quarantine

Swarmfile watches for the signature of encryption malware — a large number of *distinct entries* overwritten in a short window, scoped per-(user, machine) — and responds in two tiers. The detection mechanics are covered in [Security](#); this section is about the admin side of an incident.

- **Alert bar** (a configurable number of distinct entries in the window) — crossing it records an *alert-only* incident and notifies org owners for review. Nothing is blocked.
- **Block bar** (a higher multiple of the alert bar — twice it by default) — crossing it hard-quarantines the user: their writes stop reaching the hub (on Windows the drive stays mounted and saves are accepted locally and held; on macOS and Linux the write is refused), and lock acquire/renew both refuse, so they can't keep working even mid-lock. A single large operation can reach either bar: a 500-entry commit submitted from one machine blocks on the spot, and a 250-entry one alerts. The bars also have an aggregate backstop — one user's combined activity across their machines is judged at the block bar — so two machines doing 250 each block where one machine doing 250 alerts. If your team routinely rewrites that many files in one operation, raise the thresholds (below) rather than leaving them to be quarantined.

Owners and admins can tune either bar per organization under **Settings → Quarantine → Detection thresholds**; a blank field keeps the default, and the card shows the effective value and where it comes from. Raising the alert bar moves the inherited block bar with it (an explicit block bar doesn't move). The org's daily container-preview budget — with today's org-wide spend against it — lives on **Settings → Usage → Preview generation budgets**: previews/day default to 500 per project and 2,000 per org, a blank field keeps the default, and `0` means no limit. The source caps are fixed — video proxies accept sources up to 16 MiB (poster extraction reads the first ~16 MiB), and point-cloud (`.las` / `.laz`) previews up to 32 MiB; [File Previews](#) covers the rest.

Alert-only incidents that no one acts on auto-expire after 7 days, so a legitimate bulk edit doesn't linger as a flag (and doesn't suppress future alerts for that user). Blocks never auto-expire — they require a human to clear.

Manage incidents from the **Quarantine** view in the web dashboard. It lists active and historical incidents; each row is one click to **Release** (lift a block) or **Dismiss** (clear an alert-only flag). Admins can also manually quarantine a user out-of-band — for example, in response to an incident report that didn't trip the detector. Quarantine management is available to admins and owners.

Access revocations

Removing a member, revoking a guest, deprovisioning a user through SCIM or LDAP, revoking a device key, or invalidating a directory group membership doesn't just delete one row: Swarmfile sweeps every project in the org so cached peer credentials, wrapped keys, and access grants stop working there too. That sweep is durable — it runs across projects in the background, retries failures with backoff, and parks on a project it can't reach instead of skipping it silently. **Settings → Access revocations** (owners and admins) shows where every sweep stands: in progress, parked projects with the reason, and finished sweeps. A parked project carries a **Retry** — fix the cause and re-run the sweep from the same view; each attempt is recorded in the audit log. Revoking a project API key or a personal access token rides this same sweep to close that credential's live connections.

Because a sweep reaching a client is what makes it delete its cached copy, each engine reports a **purge receipt** when it has applied a revocation: the view shows how many devices confirmed it, how much cached data they removed, and which device each confirmation came from (attributed to the signed-in identity — a machine label is self-declared). So "dispatched" and "applied everywhere" are separate, visible facts. `GET /orgs/:org/admin/security-fanout` returns the sweep view for scripts, and `GET /orgs/:org/admin/revocation-receipts` the receipts. See [Organizations, Projects & Members](#) for member removal and [End-to-End Encryption](#) for device-key revocation.

Retention & garbage collection

Trash and soft-delete retention is configurable, from 1 to 3650 days — owners set the org-wide window in **Settings → Organization**, and an individual user can set their own window in account settings, which overrides the org default for them. The end-user side of this — restoring a deleted file, browsing history, rolling back a change — is covered in [Trash, History & Rollback](#); this page is about the org-level policy knob that governs how long deleted data is kept before garbage collection reclaims it. Set the retention window based on your org's compliance and recovery needs — a longer window costs more storage but gives you (and your users) a longer recovery runway. The default when nothing is set is 90 days.

New-project history retention. Independently of the window, every project carries a history-retention choice: **Keep every saved version** (the default) or **Named commits only**. It decides whether individual per-save versions survive the window at all — on a keep-everything project the window controls when aged history moves to cheaper cold storage, not when it is deleted, so a short window no longer means short-lived history there. Owners set the default for newly created projects with **New projects keep every saved version** in **Settings → Organization**; a creator's explicit choice in a creation form still wins, so this is a default, not an enforcement. An org whose data-minimization policy wants saves to expire should turn it off (or clear the flag on existing projects with `PATCH /orgs:orgId/projects/:projectId / swarmfile project set-keep-full-history false`).

Project storage limit

Separate from your plan's storage allowance (the size of the files themselves — see [Billing & Plans](#)), each project keeps its file information and history — names, folders, versions, locks, comments — in its own store with a fixed limit. It takes a very large project to approach it: roughly ten million files and versions combined, since a 1 TiB file costs the same here as a 1 KB one. It isn't billed.

- **Where to see it.** A project's **Metadata** tab (owners and admins) shows how much of the limit is used, whether new writes are being accepted, what the last daily cleanup removed, and — when any exist — access revocations still waiting to reach a device (the history window in force lives on the project's **Project settings** tab). **Settings → Project storage** lists every project in the org that is nearing or at its limit, worst first.
- **Early warning.** At about half the limit, the org's owners and admins get a "Project storage" notification (in the inbox, and by email unless muted per kind in Settings). History retention may also start shortening automatically — see [Trash, History & Rollback](#).
- **At the limit.** Adding files, saving new versions, copying, merging, staging, and git-LFS pushes are paused with an "insufficient storage" (`507` , `project_storage_full`) error. Nothing is lost: the Desktop App keeps changes on each person's computer and syncs them once space is available. Everything needed to recover keeps working — deleting, purging trash, restoring, renaming, moving, locks and commits.
- **Freeing space.** Purge deleted files from Trash, delete old branches, or shorten the history retention window (**Settings → Organization**, owner-only). On a keep-every-version project the window only controls cold archiving; if history genuinely isn't needed, switch the project to **Named commits only** as well. Space is reclaimed as soon as rows are removed, so writes resume without waiting for the next daily pass. A project that keeps growing past this is usually better split into several projects.

Swarmfile also watches these stores from the platform side and pages its own operators when a project approaches its ceiling, so growth is caught before it pauses writes — there is nothing for you to configure.

Long-running background jobs

Large cascades don't run inline. Above the hub's **5,000-entry** cap, folder delete, restore, purge, revert and checkout — plus `project delete` — run as durable background jobs: the request answers `202 { jobId }`, and the engine, web dashboard and tray follow the job and synthesize the original outcome (the root change lands before the `202`, so the operation's effect is already visible). Jobs are durable and survive a dropped connection, a closed laptop, or an engine restart.

From the CLI, `swarmfile op` manages them: `op list` shows recent operations; `op status <id>` reports state and live `done / total` progress (exit `0` succeeded, `1` failed, `2` canceled, `3` still running); `op wait <id>` waits and exits as the original command would; and `op cancel <id>` stops the operations that can safely stop part-way (`materialize`, `hydrate free-space`, `git backfill-oids`, `git index`, and `checkout / revert` while still staging) — a cancel that lands after the staged apply has begun is declined (`too_late`) and the operation ends with the commit's real outcome, while commits, merges and deletes complete or fail on their own and refuse (`not_cancellable`). A command that starts a long operation and doesn't follow it prints the job id and exits `3` (with `--no-wait`, or when you press Ctrl-C to detach), and the operation carries on in the engine.

Office caches

A self-hosted seed node runs headless; its operational surface is **Settings → Office caches**. The page lists the fleet with liveness (a cache reads offline after roughly 2.5 minutes without an announce, or **Awaiting approval** when your organization approves nodes), the bytes each cache served this week, and the enrollment wizard that mints a cache's credential and generates its `seed.env`. Owners and admins receive a weekly report email and one alert per day-long outage; each is muteable per-kind under **Settings → Notifications**. Setup, sizing and hardware guidance: [Self-Hosted Seed Nodes](#).

Connectivity diagnostics

Swarmfile ships two ways to run the same connectivity probe suite:

- `swarmfile doctor`, proxied through a running engine.
- `swarmfile-doctor`, a standalone binary that doesn't need a running engine at all — useful for a machine that isn't mounting anything yet, or for isolating whether a problem is the engine or the network.

Both run the same checks: DNS, hub reachability, cloud storage, the peer-to-peer layer, mDNS discovery, and OIDC/auth. Output is colorized pass/warn/fail for a human to read, or `--json` for attaching straight to a support ticket.

Diagnostics also run automatically in the background on a periodic basis, not just when you invoke them by hand, and server-side deep-health probes run from the other direction — so a problem can surface from either side. Always-public liveness endpoints exist separately, for basic uptime monitoring.

One probe worth knowing about specifically: the **plan** probe. If a self-hosted seed or NAS node is running under a plan that doesn't include the seed-node entitlement, this probe fails clearly and explicitly, instead of the node just silently misbehaving or repeating the same log line forever. If you're setting up a seed node, see [Self-Hosted Seed Nodes](#) for the plan requirement.

Reach for `swarmfile doctor` any time something feels wrong and you're not sure whether it's your network, your account, or the service — it's the fastest way to narrow that down, and its output is what support will ask for first. Full flag reference is at [swarmfile-doctor](#).

Webhooks

Webhooks let an external system react to what happens in your org — a file changed, a comment was posted, a branch was created — without polling. Swarmfile sends a signed HTTP POST to a URL you control every time something you've subscribed to happens. There are two independent kinds: an **org-wide webhook**, managed by an owner and shared by the whole org, and a **personal webhook**, a private notification channel for just your own account. Both use the same signing and delivery mechanics, described once below.

It's free and self-serve on every plan, including Free — there's no price gate and no sales conversation needed to turn it on.

How a delivery works

Every delivery is a single `POST` request with a JSON body:

```
{
  "apiVersion": 1,
  "id": "evt_...",
  "kind": "file_created",
  "data": { }
}
```

- **apiVersion** is a schema-version tag for the envelope shape itself (currently `1`). It's separate from **kind** — new event kinds get added over time without bumping this number; it only changes if the envelope's own top-level shape ever does.
- **id** is a stable, unique id for this event.
- **kind** names what happened — see [Event kinds](#) below.
- **data** is the event's own payload, shaped differently per kind.

Two headers travel with every request:

- **Swarmfile-Event** — the same value as the body's **kind**, so you can route on the header without parsing JSON first.
- **Swarmfile-Signature** — `t=<unix timestamp>,v1=<hex-encoded HMAC-SHA256>`, the same convention Stripe uses for its own outbound webhooks. Verify it before trusting anything in the body.

Verifying the signature

Compute an HMAC-SHA256 over the string `<timestamp>.<raw request body>` using your webhook's signing secret, and compare it (in constant time) to the `v1` value from the header. Use the **raw, unparsed** request body for this — not a re-serialized version of the parsed JSON, which can differ in whitespace or key order and would make the signature not match.

```
expected = hex(hmac_sha256(secret, `${timestamp}.${rawBody}`))
```

We also recommend rejecting a request if `t` is too far in the past (five minutes is a reasonable window) — this protects against a captured request being replayed later.

Node.js

```
import { createHmac, timingSafeEqual } from "node:crypto";

function verifySwarmfileSignature(secret, rawBody, signatureHeader) {
  const parts = Object.fromEntries(signatureHeader.split(",").map((kv) => kv.split("=")));
  const expected = createHmac("sha256", secret).update(`${parts.t}.${rawBody}`).digest("hex");
  const expectedBuf = Buffer.from(expected, "hex");
  const givenBuf = Buffer.from(parts.v1, "hex");
  return expectedBuf.length === givenBuf.length && timingSafeEqual(expectedBuf, givenBuf);
}

// Express example — mount with a raw-body parser so `req.body` is the
// exact bytes Swarmfile sent. Re-parsing and re-serializing the JSON
// before verifying is the single most common way to make this fail.
app.post("/webhooks/swarmfile", express.raw({ type: "application/json" })), (req, res) => {
  const signature = req.header("Swarmfile-Signature");
  if (!signature || !verifySwarmfileSignature(process.env.WEBHOOK_SECRET, req.body, signature)) {
    return res.status(401).send("invalid signature");
  }
  const event = JSON.parse(req.body.toString("utf8"));
  // ... handle event.kind / event.data
  res.status(200).end();
};
```

Python

```
import hashlib
import hmac

def verify_swarmfile_signature(secret: str, raw_body: bytes, signature_header: str) -> bool:
    parts = dict(kv.split("=", 1) for kv in signature_header.split(","))
    payload = f"{parts['t']}".encode() + raw_body
    expected = hmac.new(secret.encode(), payload, hashlib.sha256).hexdigest()
    return hmac.compare_digest(expected, parts["v1"])

# Flask example – request.data is the raw, unparsed body.
@app.route("/webhooks/swarmfile", methods=["POST"])
def swarmfile_webhook():
    signature = request.headers.get("Swarmfile-Signature", "")
    if not verify_swarmfile_signature(WEBHOOK_SECRET, request.data, signature):
        return "invalid signature", 401
    event = request.get_json()
    # ... handle event["kind"] / event["data"]
    return "", 200
```

Manual verification (curl / openssl)

Useful for confirming your receiver's math against a real delivery while debugging, without writing any code — save the raw request body Swarmfile sent to `body.json`, then:

```
TIMESTAMP="1700000000" # the "t" value from the Swarmfile-Signature header
SECRET="whsec..."    # your webhook's signing secret

printf '%s.' "$TIMESTAMP" | cat - body.json | openssl dgst -sha256 -hmac "$SECRET" -hex
```

Compare the resulting hex digest to the `v1` value from the header — they should match exactly.

Your endpoint must be reachable on the public internet. Localhost, private network ranges (`10.x` , `172.16–31.x` , `192.168.x`), and cloud metadata addresses (`169.254.169.254` and similar) are rejected at creation and re-checked immediately before every delivery, so pointing a webhook at internal infrastructure never works — this is a deliberate security boundary. Your endpoint must also respond directly: HTTP redirects are not followed.

Each delivery attempt times out after 10 seconds. Respond with any `2xx` status to acknowledge success; anything else counts as a failure for that attempt.

Org-wide webhooks

Manage these from **Settings → Webhooks**. Only owners can create, edit, or remove them — a webhook is a data-egress control, the same trust tier as an API key or an SSO configuration, not a viewing convenience open to every member.

Automating webhook management

A project-scoped API key (see [Runner \(Headless CI\)](#)) **cannot** manage webhooks, or reach any other org-wide admin surface (org settings, ACLs, directory sync) — that's deliberate: an API key is confined to one project and is never treated as an owner, no matter who minted it.

If you need to manage webhooks from a script rather than the dashboard — provisioning one as part of an infra pipeline, say — mint a **personal access token** instead, from **Settings → Access Tokens**. Unlike an API key, a personal access token isn't a separate service identity: it acts as *you*, with your own live role, so it can do anything your own account can do here and nothing more. It stops working the moment your own access does (role change, removal from the org), with no separate cleanup needed. It's a general-purpose credential, not specific to webhooks — the same token also works against any other route your role can reach.

Creating one

You need an endpoint URL. Everything else is optional:

- **Description** — a free-text label to tell webhooks apart in the list.
- **Event kinds** — leave this unset to receive every kind of event, or pick specific ones (file changes, comments, branches and tags, sharing/ACL changes, unlocks, RFIs, merge requests, uploads) to only receive what you actually care about.
- **Custom headers** — up to 5 extra `name: value` pairs sent on every delivery, alongside the standard `Swarmfile-Signature` / `Swarmfile-Event` pair. The most common use: an `Authorization` header carrying a bearer token your own receiver's gateway expects, so it can gate on that in addition to (or instead of) verifying the signature.

Once created, the signing secret is shown **exactly once** — copy it immediately. Swarmfile never displays it again; if you lose it, rotate to a new one (below).

Custom headers

A header value is exactly as sensitive as the signing secret itself — it's frequently going to be a credential — so it gets the same write-once treatment: once saved, Swarmfile can show you which header **names** are configured (so you can tell what's set without guessing), but never the values again. Editing headers means re-entering the full set you want, not patching one value in place; leaving the editor empty when editing a webhook keeps whatever's already configured, untouched.

A handful of header names are reserved and can't be overridden: `Content-Type` , `Content-Length` , `Swarmfile-Signature` , `Swarmfile-Event` , `Host` , `Connection` , `Transfer-Encoding` , and `Upgrade` — these either carry the protocol's own meaning or aren't things a webhook delivery can meaningfully override.

Delivery timing

Deliveries are dispatched on an hourly tick, not in real time — expect events to arrive within about an hour of happening, not within seconds. If your integration needs a tighter latency than that, this isn't yet the right mechanism for it.

A failed delivery is retried automatically with increasing delays (roughly 30 seconds, 2 minutes, 10 minutes, 30 minutes, then 1 hour) for up to 6 attempts total before it's given up on.

Auto-disable

If 5 events **in a row** each exhaust all of their retry attempts, the webhook is automatically disabled — Swarmfile stops trying to deliver to it rather than silently failing forever. You'll get a notification (bell, email, and your own personal webhook if you have one configured) when this happens. Fix whatever was wrong with your endpoint, then click **Re-enable** — delivery resumes from exactly where it left off, catching up on everything that happened while it was disabled, rather than skipping that window.

Editing, rotating, and removing

- **Edit** changes the URL, description, or event-kind filter without losing the webhook's id, delivery history, or audit trail.
- **Rotate secret** mints a new signing secret for an existing webhook. This is a **hard cutover** — the old secret stops verifying the instant you rotate, with no overlap window. Update your receiver's stored secret promptly; deliveries sent before you do will look unsigned to it.
- **Delete** removes the webhook and its delivery history permanently.

Testing and retrying

- **Send test** fires one synthetic event at your endpoint immediately, signed with the real secret, so you can confirm your receiver and signature verification actually work without waiting for a real event. It shows up in the delivery log as kind `webhook_test`, but never counts toward auto-disable.
- **Retry** re-sends one specific failed delivery from the log, using the exact payload that attempt originally sent. It's available for deliveries logged from now on; an earlier delivery's original payload isn't retained, since activity events themselves are derived on demand rather than stored permanently.

Both actions have a short cooldown per webhook (a few seconds) to stop a rapid click (or a script) from hammering your endpoint faster than intended.

Limits

An org can have at most 20 webhooks. Every create, edit, delete, rotate, and re-enable is recorded in the org's [Activity feed](#) under the "Security" filter, so there's an audit trail of who changed what and when.

Personal webhooks

A personal webhook is a private notification channel for your own account, alongside the bell and email — independent of any org-wide webhook an owner may have set up. Configure it under **Settings → Notifications** (the Personal webhook panel): a URL, an on/off toggle, and the same per-kind mute controls email already has (see [Notifications & Inbox](#)) — you can send comment mentions to your webhook while muting the noisier "file changed" kind, for example.

It uses the exact same envelope, headers, and signature scheme described above, with its own independently-generated secret (**Send test** and **Rotate secret** work the same way as the org-wide version).

The one real difference: a personal webhook is fire-and-forget. There's no delivery log, no automatic retry, and no auto-disable — it's held to the same bar as email, not the fuller governance treatment an org-wide webhook gets. If your endpoint is down when an event fires, that event is simply not redelivered later.

Clearing the URL turns it off. It's also cleared automatically if you're removed from the org, so a departed member's endpoint doesn't keep receiving signed events after they've lost access.

Event kinds

Event kinds match the org [Activity feed](#)'s own categories — file changes, commits, sharing/ACL changes, branches and tags, comments, unlocks, uploads, RFIs, and merge requests, plus a few webhook-governance-specific kinds. The exact, current list is what you see in the event-kind picker when creating or editing a webhook — that list is generated from the same source the Activity feed itself uses, so it can never drift out of date the way a hardcoded list on this page eventually would. A `git push` through the [git remote helper](#) counts as ordinary activity too: pushed commits, files, branches and tags land the same events a drive-side change produces, so a git workflow needs no separate hook.

Merge request events

`mr_created`, `mr_reviewed`, `mr_merged`, `mr_closed`, `mr_reopened`, `mr_assigned`, and `mr_review_requested` fire as a merge request moves through review. Every one of these carries `merge_request_id`, `mr_number`, `mr_title`, `mr_status`, `project_id`, and `by_user_id` in `data`. `mr_reviewed` additionally carries `verdict`, one of "approved", "changes_requested", or the non-blocking "commented" — so you can route on the outcome of a review without a follow-up API call. `mr_assigned` additionally carries `assigned_principal_id` and `assigned_principal_name` (the user or group that was assigned to land it); `mr_review_requested` carries the same pair under `requested_reviewer_principal_id` / `requested_reviewer_principal_name` (the user or group whose review was requested — distinct from being assigned to land it).

Every one of these also carries `assignee_names`, `reviewer_names`, and `label_names` — the merge request's *current* assignees, requested reviewers, and labels at the moment the event fired (all plain string arrays, empty if none), not just what changed. A "someone approved this" webhook consumer can tell who's on the hook for it without a follow-up API call.

A few fields are best-effort and can be absent: `source_branch_name` / `target_branch_name` are `null` if the branch itself has since been deleted past recovery, and `url` (a deep link straight to the merge request's page) is only present when the hub deployment has a site URL configured — an Enterprise on-prem deployment that hasn't set one won't include it.

`mr_created` does not fire while a merge request is a [draft](#) — on the theory that a draft isn't something you want a webhook waking anyone up for yet. That suppression is permanent, not deferred: marking a draft ready (`mr_ready`) does not retroactively fire the `mr_created` a subscriber missed at creation time, so a

draft-opened merge request may reach `merged / closed` with no `mr_created` ever delivered for it at all. `mr_assigned` and a manually-triggered `mr_review_requested` (someone explicitly requesting a review by hand) are unaffected and fire normally even on a draft — those are the events to rely on if you need a webhook signal that a specific draft exists. The one exception is `mr_review_requested` fired automatically by a branch-protection [auto-reviewer](#) at creation time: that one IS suppressed on a draft, same reasoning as `mr_created` — the reviewer is still attached, the webhook just doesn't fire for it until something later (marking the MR ready, then a subsequent manual action) generates a notifying event.

Where to go next

- [Automating Swarmfile](#) — the credentials and surfaces for driving Swarmfile from scripts and CI.
- [Notifications & Inbox](#) — the bell, email, and quiet hours, which a personal webhook sits alongside.
- [Operations](#) — the org-wide Activity feed a webhook's own governance actions appear in.
- [Permissions](#) — how ACLs shape what an org member can see and do, for context on why webhook management is owner-only.

Deploying to Your Team

For rolling Swarmfile out to a fleet rather than having each person install it by hand — silent install, GPO/MDM push, and what a freshly-provisioned machine needs (or doesn't) before someone signs in.

Windows

Prerequisites first, then the app

The download most people should use is the bundled installer (`Swarmfile-Setup.exe` / `Swarmfile-x64-Setup.exe`), which chains WinFsp, the VC++ Redistributable, and WebView2 ahead of the app itself, skipping anything already present. For a scripted mass deployment, install those prerequisites through your own tooling (GPO, SCCM, Intune) first, then push the **bare** `.msi` — it's published under its own stable name specifically for this, and unlike the bundle it does not install those prerequisites itself:

```
msiexec /i Swarmfile-x64.msi /qn /norestart
```

If a probe-backed prerequisite is missing, the MSI refuses during install with a message naming it — WinFsp or the VC++ Redistributable (instead of applying and failing later with a side-by-side or mount error). WebView2 is required at first launch but is not part of that gate; the bundle installs it, so stage it in the same task sequence when you use the bare MSI. The checks are `Installed OR ...-guarded`, so repair and upgrade are never blocked by a missing prerequisite.

`/qn` (fully silent) and `/norestart` are the same flags this project's own CI uses to verify the installer end to end, not a guess — that's a meaningfully stronger guarantee than "should work" for an MSI. Uninstall is the mirror image: `msiexec /x <ProductCode> /qn /norestart`.

The app install itself doesn't require a reboot — it closes and relaunches Explorer to pick up the shell integration, not the OS.

A scripted rollout never sees Windows' SmartScreen prompt. If someone runs a freshly downloaded `Swarmfile-Setup.exe` by hand, Windows shows **More info** → **Run anyway** until the signing identity builds up download reputation — expected, not a failed download. The `Setup.exe` / `.msi` are Authenticode-signed; the `-bin.zip` payload is not. See [Getting Started](#).

It installs per-machine, which is what makes GPO work

The MSI installs per-machine (not per-user), so a **Computer Configuration** software-installation policy is the right kind of GPO push — you don't need a per-user policy for it to reach everyone who logs into that machine. Auto-launch for the Desktop App and engine is likewise registered machine-wide, so every user account on a pushed machine gets Swarmfile running automatically at login, not just whichever account happened to run the installer. It also installs the command-line tools (`swarmfile`, `swarmfile-lfs`, `swarmfile-doctor`, `swarmfile-migrate`, `swarmfile-search`, `swarmfile-verify-history`, `swarmfile-seed`, `swarmfile-runner`, `git-remote-swarmfile`) into `C:\Program Files\Swarmfile` and appends that folder to the machine `PATH`, so a pushed machine gets the CLI with no per-user setup.

The MSI also adds an inbound Windows Firewall rule for the engine: UDP, local subnet only, Domain and Private profiles. Without it, a silently installed machine can't discover or be reached by colleagues' machines on the office LAN, and every transfer goes through the cloud. The rule is removed on uninstall. If your GPO ignores locally added firewall rules, push the equivalent rule yourself. See [Network Requirements](#) for the exact rule and a `New-NetFirewallRule` command.

First sign-in is still interactive

There's no way around this today: an employee's first launch requires them to sign in through the Desktop App. For password sign-in, that's a plain email/password form inside the Desktop App itself — no browser involved at all. For SSO, it's a real browser-based flow: the Desktop App hands off to your system's default browser for the IdP's login page via PKCE, then returns control to the Desktop App. Nothing in the product supports a fully unattended, no-user-interaction first login for a regular desktop seat, regardless of which method you use. If you need machines that never have a human sign in — a render farm node, a CI runner, a headless seed node — that's a different mechanism entirely: see [Headless and service machines](#) below.

What *is* automatic: once someone signs in, the engine adopts an org and project on its own, and it's more permissive than you might expect — it doesn't wait for an unambiguous single choice. It picks the *first* org the account can reach (even for a multi-org account — there's no picker on first run), and within that org it picks the `default` -slugged project if one exists, else just the first project in the org. Orgs start with **zero** projects: creating the first one is a deliberate action in the dashboard (or `swarmfile project create`), and until an org has one there is nothing for the drive to mount. Once every org you're rolling out has a project, "just sign in and it works" is accurate for nearly every org shape — the drive mounts to *some* project on first launch, not necessarily the one you'd have picked for them. If a newly-provisioned employee needs a specific non-default project, that's a one-click switch in the Desktop App's org/project picker after first sign-in, not something to pre-plan around. The same switch is scriptable without the Desktop App — `swarmfile workspace list`, then `swarmfile workspace switch --org <id> [--project <id>]` — and, like the picker, it needs no re-authentication: the session already covers every org the account can reach. See [Organizations, Projects & Members](#).

If your org has SSO configured with an auto-provisioning group (see [Identity](#)), that removes the manual-invite step for anyone in that IdP group — they can sign in the moment their machine has Swarmfile installed, with no separate Swarmfile-side invite to send. SCIM alone doesn't do this on its own: it syncs your directory into Swarmfile, but org access still requires that SSO group match or a manual invite.

If you're self-hosting the control plane

Everything above assumes the standard hosted service, where a fresh install already knows which hub and identity provider to talk to. A self-hosted control plane — the same software running entirely on infrastructure you operate — is available on Enterprise. With it, the deployment needs to know which hub and identity server are yours *before* first launch: pre-stage a `config.json` next to the install with your `hub_url`, `oidc_issuer`, and `oidc_client_id` set, and push that file alongside the installer through the same GPO/MDM channel. See [Engine Config File](#) for the full key reference, and [talk to us](#) to scope a self-hosted deployment.

macOS

The installer is a single signed and notarized universal `.pkg` (Intel and Apple Silicon), built to work with standard MDM package-deployment tooling (Jamf, Kandji, Mosyle, Apple Business Manager) — its own setup step recognizes when it's running under an MDM-driven install session rather than a human at the keyboard, and logs accordingly so it shows up sensibly in your MDM's policy logs rather than as a silent failure. If you hit anything MDM-specific that doesn't behave the way you'd expect, [let us know](#).

Linux

The `.deb` pulls its runtime dependencies (`libwebkit2gtk-4.1-0`, `libgtk-3-0`, `libfuse2`) in automatically through normal `apt / dpkg` dependency resolution — nothing extra to stage on a normal online install, though an offline or mirrored image needs all three available in the mirror. It enables the engine and Desktop App as per-user systemd units globally, so **new logins** pick them up with no further action; a user already logged in at install time won't have them start automatically until their next login (or you start them by hand). The Desktop App unit is conditioned on a graphical session and quietly does nothing on a headless box, which is the correct behavior for a Linux seed/NAS node.

One thing the package does *not* handle: some distributions require `/etc/fuse.conf`'s `user_allow_other` set, or the installing user added to a `fuse / plugdev` group, before a non-root user can mount via FUSE. If mounting fails with a permissions error on a specific distro, check that first — it's a FUSE/distro configuration matter, not something the `.deb` sets up for you.

The `.deb` doesn't touch your firewall either. If machines run `ufw` or `firewalld` with inbound traffic denied, allow UDP from the local subnet so LAN peers can find each other. [Network Requirements](#) has ready-made rules.

Headless and service machines

Render-farm nodes, CI runners, and self-hosted seed nodes don't have anyone signing in interactively, so they use a different credential entirely: a project-scoped API key (`swarmfile api-key create <name> --project <id>`, admin/owner-only), set as `SWARMFILE_API_KEY` on the target machine. An engine started with that env var skips the browser sign-in flow completely — no human interaction, ever — but is confined to exactly the one project the key was minted for. This is the right tool for provisioning unattended fleet machines as part of a rollout; it's a separate concern from getting Swarmfile onto employees' desktops, and the two shouldn't be conflated — a desktop seat always goes through interactive sign-in (above), a service machine never does.

Self-hosted seed nodes have a guided path aimed at exactly this: **Settings → Office caches** generates a cache's complete `seed.env` for a chosen office and project (scope, hub URL, and a freshly minted project-scoped key, shown once), then lists the whole fleet with liveness and the bytes each cache served this week. Owners and admins get a weekly report email and one alert per day-long outage, both muteable under **Settings → Notifications**. See [Self-Hosted Seed Nodes](#) for setup and sizing.

Each key can also be given its own request ceiling (`swarmfile api-key rate-limit <id> <max>`, admin/owner-only), capped at the org's per-user limit — the lever for letting a fleet burst without raising a person's window. See [Organizations, Projects & Members](#).

Repairing or reinstalling at scale

`swarmfile-doctor --repair --yes` is built for exactly this — the `--yes` flag exists specifically so it can run non-interactively from MDM/SCCM/Ansible. It still needs to run with the same elevated privileges an interactive install would (SYSTEM on Windows, root on macOS/Linux) — a scripted push job running as SYSTEM/root can drive it directly; it's not something you can trigger remotely from an unprivileged context. See [swarmfile-doctor](#).

Where to go next

- [Identity](#) — set up SSO and SCIM before you push installers, so provisioned employees can sign in and start on first launch with no separate invite to send.
- [Network Requirements](#) — the firewall allowlist to have in place before rollout.
- [Engine Config File](#) — every `config.json` key, for self-hosted or otherwise customized deployments.
- [Self-Hosted Seed Nodes](#) — turning existing on-prem hardware into a warm cache tier as part of the same rollout.

Data Portability & Offboarding

Straight answers to the questions that usually come up during procurement: how do we get our data out, what happens if we cancel or downgrade, and what does uninstalling actually remove. Where something isn't automated or self-serve, it's stated as such rather than glossed over.

Getting your data out

Because a Swarmfile project is a real mounted drive rather than a walled-off web UI, there's no special export step: copy files off the mount with Finder, Explorer, `cp`, `rsync` — whatever you'd normally use. There's no throttle or block on bulk reads while your org is in good standing, and any method below works right up until the moment you cancel — which, as [If you cancel or your payment fails](#) explains, starts a countdown rather than leaving reads available forever. If you're leaving, export first. (There is also a WAN bandwidth throttle you can optionally schedule for office hours, but that's a setting you control for your own network's sake, not a vendor-imposed export limit.)

For a scripted or continuous bulk export, [Branch Mirror](#) (Pro/Enterprise) does exactly this: it continuously exports a project branch into an S3-compatible bucket you control, as fully-hydrated real files at their real paths — not Swarmfile's internal format — a practical continuous export. The honest caveats: it's configured per branch, it's asynchronous (changes land in your bucket within a few hours — about three on an active project, about six while it is idle — not instantly), it tracks the branch continuously rather than capturing a point-in-time snapshot, so it complements a scheduled backup rather than replacing one, it's managed-tier only ([end-to-end encrypted](#) projects are refused, since the hub can't decrypt them to hydrate readable files), and it requires a Pro or Enterprise plan. One thing worth knowing in your favor: an already-running mirror keeps syncing even through the grace period described below, when the mount and web app are otherwise locked out — set it up before you need it, not after. `swarmfile-migrate`, by contrast, only moves data *into* Swarmfile, not out. If you'd rather not stand up a mirror, "mount it and copy the files" is always available as the manual export mechanism — [talk to us](#) if you need help.

If you use Swarmfile as a **git-LFS server**, there's a third exit that's fully self-serve and needs no Swarmfile software at all: `git lfs fetch --all` pulls every LFS object into your local git-LFS store with stock git-LFS, `git lfs fsck` verifies each one, and you then repoint `.lfsconfig` at any other LFS server (GitHub, GitLab, an S3-backed server, self-hosted) and `git lfs push`. Because your source lives in your own git host and only the binary objects are in Swarmfile, this gets those objects out with standard tooling and no vendor dependency — see [Leaving Swarmfile](#) in the git-LFS guide. On each machine that pushed large objects, run `swarmfile-lfs uninstall` to remove the transfer-agent wiring from git config once you're done (committed `.lfsconfig` files and stored credentials are separate — the command tells you what it left alone).

If you cancel or your payment fails

Both paths lead to the same place: a **28-day grace period** during which the org's content is inaccessible, followed by permanent deletion. There is no indefinite post-cancellation read state — treat cancellation as a deletion deadline and get your data out first (see [Getting your data out](#)).

If you cancel deliberately (via the Stripe Customer Portal): the org leaves its paid subscription state (there's no free tier it falls back to — it doesn't become a Free-plan org), and for the first ~24 hours nothing else changes, a deliberate cushion so an accidental click can be undone by resubscribing. After that window the org is flagged as unsubscribed and the standard grace period begins, exactly as if a payment had failed:

- **A real 28-day grace period runs, tracked independently of Stripe's own retry schedule** — from the moment the cancellation is flagged, not shortened or extended by whatever dunning timeline Stripe itself runs.
- **Reads are blocked too, not just new uploads — everywhere, not only the mount.** The desktop mount, the web dashboard's own download/preview, and public share links you've handed out all stop serving that project's content for the whole window, not only once the deadline nears. The dashboard and billing page themselves stay reachable so you can resubscribe. The one deliberate exception is [Branch Mirror](#): an already-running mirror keeps exporting through the grace period, covered above.
- **We email the org owner as the window closes** (day 1, 7, 21, and the 27th day) so there's real warning before anything happens, not a surprise.
- **If the 28 days pass with nothing resolved, the org's data is permanently deleted — and so is the owner's account, if they don't belong to any other org.** This includes every project, file, and version in the org, the underlying storage, and the org itself; it also includes the owner's Swarmfile login if that org was their only one (an owner who's a member of another org keeps their account either way). This is genuinely irreversible.

If a payment fails and is never resolved (a declined renewal, an expired card, and Stripe's own retries are exhausted): the same 28-day window starts from the moment the payment first failed, with the same blocked reads, the same reminder emails, and the same permanent deletion at the end.

Resubscribing at any point before the 28 days are up cancels the countdown entirely and restores access, with nothing lost. If a payment is failing and you want your data out, do it before the window closes — and if you want to leave without a deadline, either export first and then cancel, or ask us to delete the org outright (below) once your export is done.

If you downgrade over your new plan's limit

Downgrading (Pro to Starter, for instance) when your stored data exceeds the smaller plan's included allowance doesn't delete or lock anything either. The included allowance is a billing threshold, not a hard ceiling — you'd simply be billed overage on the excess, exactly as if you'd grown into it gradually on your current plan. New uploads only actually get blocked if you cross a much higher hard ceiling (an abuse backstop, well above what any plan advertises as included) — existing files remain fully readable at every point in between.

Deleting an organization entirely

If you want an org torn down on your own schedule — not because of a lapsed payment — the org owner can do it from the web dashboard: **Settings → Organization → Delete organization**. It asks you to retype the org slug, and to enter your password if your account has one; it must be done from an interactive sign-in — an API key or personal access token can't perform it. It then cancels any active subscription, revokes the org's API keys, disconnects any connected GitHub App installation, and deletes the org, its projects, its files, and its stored data. The purge writes durable audit rows that outlive it, so the org deletion remains auditable after the data is gone; if the credential revoke stalls, the purge reports `credential_revoke_pending` rather than success. It is immediate and irreversible, so make sure your export is complete first. Admins and members won't see the control; it's owner-only. If you'd rather have us handle the deletion, [contact us](#) and we'll do it directly.

The automatic 28-day path described above is the one exception: that deletion is triggered by billing status, not a support request, for the reasons explained there.

Uninstalling the desktop client

Uninstalling a machine never touches your account or org data — that all lives on the hub, not the machine you're uninstalling from. What the uninstaller does *not* do, on any platform, is clear that machine's local cache (the block cache, upload queue, and local metadata database) — that's left in place, in case you're reinstalling rather than fully decommissioning the machine.

The one real risk: Swarmfile's writes are designed to return instantly and upload in the background — a save is only visible to your teammates once it's actually finished uploading to the hub, which for a large file can take a while after the save itself completes locally. If you uninstall (or, more precisely, wipe that machine's local Swarmfile cache) while an upload is still queued and hasn't finished confirming with the hub, that specific save is lost — it never makes it to the hub, and nobody else ever sees it. Reinstalling on top of the same, un-wiped cache resumes the queue and finishes the upload normally; wiping the cache and starting fresh does not. In practice: before decommissioning a machine, check the Desktop App for pending uploads and let them finish first.

Self-hosted seed nodes as an exit ramp

A [self-hosted seed node](#) fetches and mirrors a project's data through hardware you own — keeping a warm local cache on it — which sounds like a natural answer to "what if the hosted service goes away." For the standard managed (encrypted) tier — the default, and what most projects use — it's a more qualified answer than that framing suggests: what the seed node holds is ciphertext, and reading it depends on a decryption key fetched live from the hosted hub. There's no offline fallback for that key. So a seed node is a genuine speed and reliability win for your office day to day — most reads served locally, less dependence on your internet connection — but it isn't an independently readable backup you could fall back to if the hosted service were ever permanently unreachable.

If true offline independence — not just a warm local cache, but a fully self-contained deployment with no dependency on our hosted infrastructure at all — is a hard requirement, that's the Enterprise self-hosted control plane: the same code running entirely on infrastructure you control, available on Enterprise (a fully air-gapped deployment is on the roadmap) — see [Deployment Topologies](#) and [talk to us](#) about what it would involve for your environment.

Where to go next

- [Billing & Plans](#) — how plan switching, seats, and overage actually work.
- [Self-Hosted Seed Nodes](#) — running your own warm cache tier.
- [Security Architecture](#) — the full key-custody and threat-model detail behind the encryption note above.

Branch Mirror

Branch mirror continuously exports one project branch into an S3-compatible bucket you control — as **real files at their real paths**, not Swarmfile's internal content-addressed block format. It's a secondary, one-directional copy: our storage platform stays the primary, authoritative home for your data, and the mirror is a fully-hydrated, browsable duplicate that lives in your own bucket for backup, disaster-recovery, or downstream-pipeline use.

Available on **Pro and Enterprise** (the `s3_mirror` entitlement), configured per branch from the dashboard.

How it's different from the alternatives

- **Not `swarmfile-migrate`**. That tool is a one-shot import that moves data *into* Swarmfile (from a NAS, S3 bucket, or file list). A branch mirror runs the other direction and keeps running — a continuous export *out* to a bucket you own, not a single copy-in.
- **Not Dedicated Storage**. **Dedicated Storage** is about isolating your bytes *inside our own infrastructure* — a bucket that's physically yours but still one we create and operate. A branch mirror puts a readable copy in *your* infrastructure entirely, on a provider and account you control.
- **Not bring-your-own-storage**. **Bring-your-own-storage** makes your bucket the *primary* home where the live blocks actually live. A branch mirror is strictly a read-only export copy — the authoritative data still lives on our platform, and the mirror just tracks it.

The key thing a mirror gives you that the block store can't: the files land in your bucket at their natural paths (`projectroot/scenes/shot_010.exr` , not an opaque content hash), directly readable by anything that speaks S3 — no Swarmfile client, no decryption step, no reassembly.

Where to configure it

Open the project in the dashboard and go to its **Mirror** tab. Configuration is **per branch** — each branch of a project mirrors independently (or not at all), so you can export just `main` while leaving working branches unmirrored.

The tab is **owner/admin-gated** and requires a Pro or Enterprise plan. On a plan without the `s3_mirror` entitlement the tab explains the gate; an existing mirror keeps its last-synced state and can still be turned off, but its configuration can't be changed until the plan includes mirroring again.

Configuration fields

Field	Notes
Endpoint URL	Your provider's S3-compatible endpoint. Not needed for AWS S3 itself if you use the region-default endpoint; required for Wasabi, Backblaze B2, MinIO, and most others.
Region	The bucket's region.
Bucket name	The target bucket. It must already exist — Swarmfile writes into it, it doesn't create it.
Key prefix	Optional. A sub-path inside the bucket to confine the mirror to (e.g. <code>swarmfile/</code>). Leave blank to write at the bucket root.
Access key ID	The S3 access key ID for a credential with write access to the bucket.
Secret access key	Stored encrypted at rest (AES-256) and never returned by the dashboard. On an existing mirror, leave it blank to keep the stored credential — you only re-enter it to change it.

The credential is validated with a real round-trip against your bucket at save time, so a bad endpoint, region, bucket name, or key is caught immediately in the settings screen rather than silently on the first sync.

Behavior and limits

- **Full-fidelity tracking**. The mirror reflects the branch as it changes: new and modified files are re-exported, and **deletes, restores, and renames propagate** to the target bucket. What's in the mirror tracks what's on the branch.
- **Asynchronous, up to a few hours of lag**. Syncing is driven by a periodic maintenance scan — about every three hours while the project is active, backing off to about six hours while it is idle. So a change can take a few hours to appear in your bucket (up to ~3 h on an active project, ~6 h on an idle one). This is not a live, write-through copy.
- **Fire-and-forget**. There's no delivery guarantee, no built-in retry-until-success alerting, and no per-file receipt. A file that fails to sync on one pass is simply left unconfirmed and picked up again on the next scan's fresh diff — the sync self-heals over time rather than escalating a failure to you.
- **One oversized file skips; a bad credential fails the whole sync**. If a single file trips a size cap (or hits a transient error), that one file is left unsynced and retried next tick — the rest of the batch still goes through. A broken or rotated-away credential, by contrast, fails every file, so the whole sync fails until you fix the credential.
- **Status is visible in the tab**. The Mirror tab shows the current state — **Never synced**, **Syncing**, **Synced**, or **Sync failed** — along with the last sync time and, on failure, the last error.
- **Turning it off leaves the bucket alone**. Disabling a mirror stops future syncs but never deletes what's already been written — that data is yours in your bucket.

End-to-end encrypted projects are not eligible

A mirror is **refused at configuration time for end-to-end (customer-managed) encrypted projects**. The hub never holds an E2E project's key, so it structurally cannot decrypt the blocks to hydrate readable files — there's nothing it could write to your bucket but ciphertext. Mirroring is a managed-tier feature; on an E2E project the save is rejected with an explanation rather than producing an unusable export.

Data residency is your bucket's responsibility

Once a file is written to your bucket, where it physically lives is governed by *your* provider, account, and region — not by Swarmfile's [data residency](#) pinning. Swarmfile can't verify or enforce the residency of an arbitrary customer-supplied endpoint, so for an org with a **pinned jurisdiction** — or for any individual **project pinned to its own region** — a branch mirror is refused outright rather than silently allowed to route data outside that jurisdiction. If you're using a mirror, choosing a bucket whose region matches your own residency commitments is on you.

Where to go next

- [Bring-your-own-storage](#) — making your bucket the *primary* home for live blocks (Enterprise).
- [Dedicated Storage Isolation](#) — physical isolation inside our own infrastructure.
- [Data Portability & Offboarding](#) — the fuller picture on getting your data out.

Filesystem Compatibility & Conformance

Swarmfile mounts as a real drive — `~/Swarmfile` on macOS and Linux, a drive letter on Windows — through FUSE/FUSE-T and WinFsp, not through a synced local folder. The exact location is per-machine and can be configured (the engine's `mount_point` ; Windows picks its drive letter at runtime), so read yours from the Desktop App's Status pill rather than hard-coding a path. That distinction matters for exactly the kind of file your native apps care about: does a rename actually rename, does a lock actually stop a second writer, does opening a 200 GB file touch the whole thing or just the bytes you read.

This page states plainly what works today, per platform, and what's a known gap.

What's verified

Most rows below are exercised by automated write-semantics tests that drive the *mounted filesystem* the way an application does — not the internal storage code directly — on a live mount per platform. A few rows (adapter capabilities and platform-specific behavior) are covered by targeted tests instead; the Notes column says when.

Capability	macOS	Linux	Windows	Notes
Create / read / write	Verified	Verified	Verified	
Delete (file & directory)	Verified	Verified	Verified	On Windows, verified via Win32 <code>RemoveDirectory</code> , <code>cmd rmdir /s</code> and .NET's recursive <code>Directory.Delete</code> , including on empty folders. PowerShell's <code>Remove-Item -Recurse</code> fails through the mount (known limitation — use one of the above instead)
Rename — same directory	Verified	Verified	Verified	
Rename — cross-directory (move)	Asserted	Asserted	Verified	Entry identity, history, comments, and locks move with the file — this is not a copy-then-delete. Windows verified on a live mount; macOS/Linux covered by the conformance suite
Exclusive create (<code>O_EXCL</code> / Win32 <code>CREATE_NEW</code>)	Verified	Verified	Verified	A create that races an existing name is refused (<code>EEXIST</code> / <code>STATUS_OBJECT_NAME_COLLISION</code>), enforced at our layer — so it holds even offline, not just when the OS routes the open
Overwrite / truncate an existing file	Verified	Verified	Verified	Content-checksummed, not just size-checked
Directory listing at scale	Verified	Verified	Verified	Wildcard/exact-name filters are verified on Windows only — POSIX shells and apps glob client-side over a plain readdir, so there's no equivalent operation to assert there
Case handling	Case-preserving	Case-preserving	Case-insensitive	Matches each platform's native convention. Names also resolve across Unicode forms on every platform: <code>café.txt</code> created as NFC opens when asked for the NFD spelling, an ASCII-case-only difference resolves, and the Turkish dotted/dotless I family (<code>İ</code> / <code>i</code> vs <code>ı</code>) is treated as one name — the store keeps the spelling you created with. Verified live on Windows in both request directions and unit-tested cross-platform; the WinFsp boundary reports the requested spelling when the stored one differs by more than case, which is what lets the open succeed
Free-space reporting (<code>statfs</code>)	Verified	Verified	Verified	
<code>fsync</code> durability	Verified	Verified	Verified	A write is not reported durable until it actually is; on Windows this is the <code>FlushFileBuffers</code> path
Extended attributes (get/set/list/remove)	Local-only, via macOS's <code>xattr</code> files (see below)	Implemented	Not implemented	Windows EA specifically — see below. Not the same thing as alternate data streams, the row below. On Linux the adapter implements all four operations; nothing on a Linux mount creates them the way macOS does, so end-to-end use isn't exercised
Alternate data streams — create/read/write/enumerate/delete	N/A	N/A	Verified, incl. sync between two clients	Windows-only concept (<code>file.ext:streamname</code>). A stream also survives a copy through the mount rather than being silently dropped, the way most non-native sync tools drop it. The <code>Zone.Identifier</code> stream (Windows' Mark of the Web) stays on the machine that wrote it and never syncs — see Sync Exclusions . Every other stream's content travels the same hub upload/commit pipeline as file content. Verified on real Windows mounts: these operations, and sync between two clients through the hub. A stream written or changed on one client reads back identical on the other, a stream deletion reaches the other client, and a <code>Zone.Identifier</code> never leaves the client that wrote it
Alternate data streams — rename	N/A	N/A	Not supported (see below)	
Symlinks — create	Verified	Verified	Not supported (see below)	
Symlinks — read / follow / list / delete	Verified	Verified	Read, follow & list	On Windows, <code>readlink</code> (reparse-point resolution), following a link to its target, and listing it as a reparse point are all verified; delete-through-the-mount isn't verified on Windows
Executable bit (<code>chmod +x</code> / <code>-x</code>)	Verified	Verified	Preserved (no executable bit)	Any <code>x</code> bit sets the project's per-file executable flag; the drive reports <code>0755</code> / <code>0644</code> . A create with mode <code>0755</code> (including <code>git checkout</code> of a <code>100755</code> file) comes out executable. Windows has no mode bit: edits and safe-saves there keep the flag, and <code>swarmfile chmod</code> sets it from any platform
Byte-range locking (API)	Verified	Verified	Verified	Hub-enforced acquire/release/overlap-detection, driven over the engine control socket — the same primitive a native worksharing plugin (e.g. for Revit) would call. Release is always an explicit unlock call over that socket; on Windows in particular, a lock is not auto-released when the OS closes the file handle, so a caller must unlock explicitly. Revit/SolidWorks worksharing itself needs no plugin on Windows — Swarmfile honors the native OS file lock those apps already take and turns it into a hub-enforced claim across machines; on macOS/Linux only cooperative locks apply. A deployment that needs an open refused whenever the hub can't be reached can enable that (see Environment Variables). This byte-range API is the lower-level primitive for finer, element-level integrations beyond that. See Swarmfile vs. LucidLink for how the underlying lock guarantee compares
Folder/file ACL enforcement	Verified	Verified	Verified, incl. Explorer-visible permissions	Enforced hub-side for every client; only Windows additionally enforces and presents ACLs at the mount itself (Explorer-visible DACL). On macOS and Linux, Finder and <code>ls</code> show the mode bits from <code>getattr</code> , not the project's ACLs

Declared gaps

Each of these is a documented limitation:

- **Hard links.** Not supported on any platform. Our metadata model is one entry = one path = one content id; a hard link needs a content-identity indirection the schema doesn't have.
- **Symlink creation on Windows.** Windows mounts read and follow symlinks correctly but refuse to create one. Creating a Windows symlink as an unprivileged process requires Developer Mode or a specific privilege — support would work for some users and silently fail for others on the same drive, so we refuse

consistently instead.

- **Renaming an alternate data stream itself.** Renaming the base file that carries a stream works normally, and the stream travels with it. Renaming just the stream — `file.txt:secret` to `file.txt:hidden` — is out of scope and refused outright, not silently mishandled: we won't rename the base file when you only asked to rename one of its streams, and we won't leave the stream qualifier in a request that no path would ever resolve to.
- **Extended attributes on Windows.** Not implemented — a deliberate scope decision, separate from alternate data streams (see the verified-capabilities table above, which streams *are* supported).
- **A save the hub refuses reaches the app as `EIO`.** When a save is refused (a file held elsewhere, a plan or policy refusal, a root-create restriction, quarantine), the application only sees a generic I/O error; the reason is written under **Why did my save fail?** in the tray and printed by `swarmfile sync stuck`. The errno is deliberately generic, so read the reason there rather than from the app's error dialog.
- **`fallocate` / `copy_file_range` / `lseek(SEEK_HOLE)`.** Not implemented on the POSIX adapters.
- **`fcntl()` byte-range advisory locks are kernel-local only.** They don't coordinate across machines — that's a different mechanism from the hub-enforced byte-range locking API above. An app that relies on plain POSIX advisory locks for multi-writer coordination (rather than calling into the lock API) gets locking that looks like it works but only protects against other processes on the same machine.
- **Windows: a memory-mapped view can outlive the handle for reads, not for writes.** Reading through a memory-mapped view works, including when an application closes the file handle before reading the map — `git` does exactly that with `.git/config`, which is why `git init/commit/status` work on the mount. Writing through a mapping is only served while a file handle is open: an application that maps a file **read-write**, closes the handle, and then writes through the view can have those write-backs fail. Keep the handle open until the mapping is unmapped, or write through the handle. Read-only mappings (the common case, `git` included) are unaffected.
- **A file isn't readable the instant it's written — for about a second, on Windows.** The engine itself publishes the new content *synchronously* as the flush completes (the local metadata row, size included, is updated before close returns), so a script that opens the file through a fresh handle gets the new bytes on macOS and Linux. Windows can still report the old attributes (0 bytes for a new file, the previous size for an overwrite) for up to its **1-second Cache Manager revalidation timeout**, because the kernel caches file info between revalidations. It matters for any workflow that writes a file and immediately reads it back — a save-then-verify step, a render job that re-opens its own output, or an app that reloads the file right after saving. A person clicking around is usually slower than the one-second window, but not always: reopening a just-saved file within that second can still see the old size. If you need certainty, wait on the engine rather than the clock: `swarmfile uploads --wait` blocks until the upload queue is empty, and `swarmfile status` splits the rest — `pending uploads` (blocks still going up), `commits pending` (bytes up, commit not landed; a plan-held commit shows under `quota-held commits`). (A plain `sync` forces pending writes to flush locally but does not wait for the hub upload and commit, so use the counters, not `sync`, when that distinction matters.)
- **`access()` / `faccessat()` aren't enforced locally.** Permission is decided hub-side by the ACL system, not by local file mode bits, so a pre-flight `access()` check reports success even for a write the hub will later refuse.
- **A read can WAIT, for a file that is still uploading.** A file that has never finished uploading can be opened — it reports the size it is going to be, and reads work. A read into a range that has not arrived yet waits for it, bounded at 15 s by default, and returns a *retryable* error — `EAGAIN` on macOS and Linux, `STATUS_IO_TIMEOUT` on Windows — rather than an I/O error if the wait times out, because an application that sees an I/O error part way through a file usually decides the file is damaged and discards your document. A stall is the correct answer for a player; a corrupt-file error is not. It only ever applies to a file with no committed version yet, so nothing you are already reading changes behavior.
- **macOS: `ditto` works.** Plain `cp / cp -R`, `rsync`, Finder, and the installers and build scripts that reach for `ditto` all work.
- **macOS extended attributes are local-only.** The drive is mounted without NFS named attributes, so macOS keeps Finder tags, comments, and the like the way it does on any such volume: in an AppleDouble `._name` file beside the original. Swarmfile keeps those `._` files, and `.DS_Store`, on the machine that wrote them and **never syncs** them (see [Sync Exclusions](#)) — `cp -p`, `ditto`, and `xattr` work with them on that machine, but a teammate (or that machine after a fresh install) won't see them. Windows alternate data streams are the exception and do travel (except the `Zone.Identifier` Mark of the Web, which stays local); POSIX `xattrs` crossing machines would be a hub-side format addition, not a mount fix.

How we verify this

Each platform has its own conformance suite driving the live mount, and a shared manifest keeps the three platforms' coverage aligned — a capability missing on one platform is caught rather than left for you to find. Known gaps are listed above.

How this compares to typical cloud storage clients

Most general-purpose cloud storage clients — the kind built primarily for syncing documents and photos — are sync engines wearing a filesystem's clothing: a background process reconciling a local folder against the cloud, not a filesystem driver with defined semantics for the operations above. That's a fine model for a folder of PDFs. It tends to fall over for exactly the workloads Swarmfile targets: whole-file download before an app can open anything, no cross-process byte-range locking, and "two people edited it" resolved by keeping both files rather than by a lock.

This isn't a hypothetical concern — it shows up in vendor documentation for the professional tools this product is built for. Autodesk's own support article on [using cloud-synced services with Revit files](#) states that file-based worksharing isn't supported on common cloud-sync storage, and names corruption of the central model and lost work as the consequence. Esri's knowledge base similarly [documents common problems running ArcGIS Pro against cloud storage services](#). If you're evaluating any filesystem — including this one — for this kind of workload, the questions worth asking are the ones this page answers: is locking enforced per byte range or just per whole file, is it enforced on every platform your team actually uses, and is any of it independently verified rather than asserted.

Supported Platforms & System Requirements

Swarmfile ships desktop installers for macOS, Windows, and Linux, plus command-line tools with each. This page lists what runs where, what each installer includes, and the limits worth knowing before you plan a deployment.

At a glance

Platform	Architectures	Installer	Minimum OS
macOS	Universal (Apple Silicon + Intel)	Swarmfile.pkg (full install), Swarmfile.dmg (app only — no service or PATH setup)	macOS 11.0
Windows	x64, ARM64	Swarmfile-Setup.exe / Swarmfile-x64-Setup.exe (bundled), Swarmfile.msi / Swarmfile-x64.msi (bare, for IT)	Windows 10 or Windows 11
Linux	x86_64 only	swarmfile-amd64.deb	Ubuntu 22.04+ or an equivalent current Debian-based distro

Download them from swarmfile.com/downloads, or ask for the stable releases.swarmfile.com/latest/... URLs for a scripted install.

macOS

- **One universal build** for Apple Silicon and Intel; **macOS 11.0 or later**, enforced by the installer.
- Use the `.pkg`: it installs the Desktop App, the engine as a LaunchAgent, the macOS mount helper, and symlinks for every CLI tool into `/usr/local/bin`.
- The `.dmg` installs the Desktop App, whose bundle carries the engine, the runner, and the CLI binaries — but it doesn't set up the engine's LaunchAgent service, the mount helper, or the `/usr/local/bin` symlinks, so the engine isn't running as a service and the CLI isn't on `PATH`. On first launch the tray opens **Diagnostics** with the Reinstall action so one click puts them in place.
- Mounting is **kext-free**: Swarmfile bundles FUSE-T, so there is no kernel extension, no Recovery Mode step, and no system-extension approval on Apple Silicon.
- **macOS 15+** additionally requires **Local Network** permission for LAN discovery; the app asks for it, and headless engines running outside the app bundle need to inherit it (or run as a LaunchDaemon). See [Network Requirements](#).
- Installers are signed with a Developer ID and notarized/stapled.

Windows

- **x64 and ARM64** builds; Windows 10 or Windows 11. Both install **per-machine** into `C:\Program Files\Swarmfile` and add that folder to the machine `PATH`, so the CLI is available to every account.
- **Use the bundled Setup.exe for people**: it chains WinFsp, the Visual C++ Redistributable, and the WebView2 runtime, skipping whatever is already present. Windows Server/VDI images usually need those three staged by your own tooling first.
- **Use the bare .msi for scripted rollouts** (`msiexec /i Swarmfile-x64.msi /qn /norestart`): it installs only the app and expects those prerequisites to be present. It checks for WinFsp and the VC++ Redistributable during install and refuses with a message naming whichever is missing, rather than failing later with a side-by-side error; WebView2 isn't part of that gate, so stage it too. Repair and upgrade are never blocked by these checks. See [Deploying to Your Team](#).
- One Windows-specific behavior worth knowing: **symlink creation is refused on the mount**. Windows needs Developer Mode (or a specific privilege) for an unprivileged process to create symlinks, so Swarmfile refuses consistently instead of failing for some users; reading, following, and listing symlinks all work.
- The `Setup.exe` and `.msi` installers are Authenticode-signed; the `.zip` tools and the standalone `swarmfile-doctor` download are not, so Windows shows the SmartScreen prompt on first run — see [Getting Started](#).
- The bundled **WinFsp is pinned to a tested 2.2 Beta build** (the exact pin is carried in the Windows installer's arch map); the pin moves to 2.2 stable when it ships.

Linux

- **x86_64 only**. The `.deb` requires `libwebkit2gtk-4.1-0`, `libgtk-3-0`, and `libfuse2`; install it with `apt` and the dependencies resolve normally.
- It installs **per-user systemd units** (`swarmfile-engine`, `swarmfile-tray`) enabled globally, so new logins start them with no extra action; a user already logged in needs to log out and back in (or start the units) for them to take effect. The Desktop App unit is conditioned on a graphical session, so it does nothing on a headless box.
- The `.deb` **does not touch your firewall**; on a host with a default-deny inbound policy, allow UDP from the local subnet — see [Linux firewalls](#).
- **.rpm isn't published**. If you run RHEL/Fedora/OpenSUSE, [talk to us](#) about a supported `.rpm` path.

Hardware and disk

There is no formal RAM or CPU minimum: the engine is a native service, and the mount is a filesystem driver, both modest by desktop standards. The numbers that matter are disk-related:

- **Local cache: 10 GiB by default** (`cache_max_bytes`). Files you touch are cached locally and evicted as the cache fills; it is not a copy of the project. Adjust it in the Desktop App's settings or `swarmfile cache set <GiB>`.
- **The mount reports at least 1 GiB of free space** to applications even when the cache is full but evictable, so a "disk full" error from your app reflects local disk pressure, not the project's size.
- **A seed/NAS node** is different: it fetches whole projects (mirroring them to cloud storage) and keeps them in a warm cache with its own budget. Size `SWARMFILE_CACHE_MAX_BYTES` for the working set you want resident — the sample seed config suggests a **100 GiB cache as a starting point** and raising it to 1 TiB or more on a capable NAS; an evicted block re-fetches on the next read.

- **Project metadata**, not file bytes, has a per-project limit of roughly **ten million files and versions combined** — very hard to reach unless you keep every version of a very large tree forever. See [Project storage limit](#).

File and object size limits

Content is chunked (FastCDC, 256 KiB–4 MiB chunks) and manifests are segmented once they grow past ~40,000 chunks, so a **file on the mounted drive has no practical size cap** — multi-terabyte files work, and the manifest structure extends rather than truncating. Two specific ceilings are worth planning around:

- **git-LFS objects: 256 GiB maximum.** A single LFS object beyond that is refused. Stock `git lfs push` works for objects up to about 5 GiB; larger pushes need the desktop transfer agent — see [git-LFS](#).
- **Hub-mediated downloads** assemble up to 900 manifest segments per read; segmented manifests keep large files readable. A very large *native* `git-lfs` download may first be paced while the hub prepares it (the client's retry setting carries it through), or refused on a deployment without the prepare queue — the mount is the read path there.

Platform capability differences

The filesystem semantics are the same everywhere for ordinary work (create, read, write, rename, delete, locking). A short list differs by platform:

Capability	macOS	Linux	Windows
Symlink creation	Yes	Yes	Refused (reading/listing works)
Extended attributes	Kept locally, not synced (<code>cp -p / ditto</code> work)	Adapter implemented (unit-tested; local-only)	N/A — use streams (streams do sync)
Alternate data streams	N/A	N/A	Yes (rename of a stream itself is refused)
Filtered/wildcard directory listing	App globs client-side	App globs client-side	Native filters verified
ACLs at the mount	Hub-enforced; <code>ls</code> shows mode bits	Hub-enforced; <code>ls</code> shows mode bits	Hub-enforced and projected as Explorer DACLS
<code>ditto</code>	Works	N/A	N/A
PowerShell <code>Remove-Item -Recurse</code> on a directory	N/A	N/A	Not supported — the cmdlet fails through the mount; use <code>cmd rmdir /s</code> , .NET's <code>Directory.Delete</code> , or Explorer

The full per-operation table, including the declared gaps (hard links, `fcntl` advisory locks, sparse files), is in [Filesystem Compatibility & Conformance](#).

Where to go next

- [Release Channels & Updates](#) — how installs update and reinstall.
- [Deploying to Your Team](#) — silent install and MDM/GPO rollout.
- [Network Requirements](#) — the firewall allowlist before rollout.

Release Channels & Updates

Swarmfile installs update themselves, but never silently and never without you saying so. This page covers where installers come from, how updates are delivered and verified, and how to reinstall a machine cleanly.

Where installs come from

What	Host
Installers	<code>releases.swarmfile.com</code>
Update manifest	<code>updates.swarmfile.com/stable/latest.json</code>
Dashboard / web app	<code>swarmfile.com</code>
Hub / identity	<code>hub.swarmfile.com</code> , <code>id.swarmfile.com</code>
git-LFS agent installer	<code>get.swarmfile.com/lfs</code>

A machine's hosts are baked into its install, not compiled into the binary: packaging writes the hub, identity, site, and update endpoints into the bundled `config.json`, so an install keeps the endpoints it was packaged with. See [Engine Config File](#).

Installers are available two ways under `releases.swarmfile.com`:

- **Stable aliases** — `releases.swarmfile.com/latest/Swarmfile.pkg`, `.../latest/Swarmfile.dmg`, `.../latest/swarmfile-amd64.deb`, `.../latest/Swarmfile-Setup.exe` (Windows ARM64), `.../latest/Swarmfile-x64-Setup.exe`, and the bare `.msi` files. These are what the download page's links resolve to (the page links through the hub, which redirects here).
- **Versioned paths** — `releases.swarmfile.com/v<version>/Swarmfile-<version>.pkg` and so on, so a scripted install can pin an exact build whose SHA-256 stays valid. Each artifact ships with a `.sha256` file and there is a combined `SHA256SUMS` per release. Homebrew, winget, and Chocolatey packages point at the same signed artifacts.

Use the **stable aliases** for people and the **versioned paths** for automated mass installs.

How updates work

Updates live in the Desktop App, not the engine:

- It checks once at startup and then **every 45 minutes** while running. Repeated failures back off, up to about every six hours.
- A found update **only** raises a notification: a pill in the header and a card in **Diagnostics**. Swarmfile **never installs an update on its own** — **Install now** is the only path, and that is deliberate: an unattended restart while a save or render is running is worse than being a few days behind.
- If an update sits uninstalled for **a week**, the Desktop App sends **one** desktop notification naming the version — still informational, still install-on-your-command. (A restart re-raises the pill, so the nudge is for the tray that stays running for weeks.) You can also press the manual **Check for updates** action in Diagnostics at any time.

What happens when you press **Install now** differs by platform:

- **macOS and Windows** — the app downloads the signed package and installs it, then restarts itself (and its engine). On Windows an unelevated app starts the **Swarmfile Updater** helper (an on-demand, never-auto-triggering elevated helper) which verifies the package signature and applies it.
- **Linux** — the app cannot self-install (it runs as your user, and package installation needs root). It shows a notification pointing at the `.deb` for you to install; there is no silent privilege escalation.

There is no beta, nightly, or insider channel. Every release publishes to the single `stable` channel.

Update integrity

Updates are verified before they are applied:

- The update manifest's artifacts are signed with a project **Ed25519 (minisign)** key. The public key is compiled into the app and into the Windows updater helper, so a compromised update server cannot hand you a package that verifies — the endpoint can be moved, the trusted key cannot.
- The **app itself** refuses an update whose signature does not verify, and refuses a manifest that is not strictly newer than the installed version.
- **macOS** installers are Developer ID-signed and notarized/stapled. **Windows** installers (`Setup.exe` and the bare `.msi`) are Authenticode-signed in the release pipeline; the `-bin.zip` payload and the standalone `swarmfile-doctor` binary are not signed. **Linux** `.deb` s are not distro-signed (no apt repository key); verify them against the published `.sha256 / SHA256SUMS` and rely on the updater's own signature when updating. On Windows, the SmartScreen prompt is expected on first run either way — see [Getting Started](#).
- Installing an **older** version over a newer one is not supported: the Windows MSI refuses downgrades, and the updater refuses non-newer manifests. To move a machine backward, remove the app and install the exact older package from its versioned path.

Version numbering

Versions are `0.<minor>.<patch>`; the current line is `0.3.x`, and the newest published version is whatever the stable release manifest reports (`updates.swarmfile.com/stable/latest.json`). The version is embedded in the installer name and in every versioned URL, and release publishing is monotonic — an older version can never overwrite the current `latest` manifest. The engine, CLI tools, and Desktop App in one install share the same version; if they ever drift — an update that leaves a background service pointing at an older engine binary, say — the app raises a version-mismatch banner and offers the right remedy: restart or install for a plain mismatch, and repair/reinstall when the service itself points at the old binary, where a restart would just re-run it.

Reinstalling or repairing

When an install is unhealthy — a broken mount helper, a missing CLI tool, a half-applied update — reinstall rather than removing things by hand:

- **Desktop App: Diagnostics → Reinstall...** (Windows and macOS).
- **Command line, headless, or scripted:** `swarmfile-doctor --repair --yes`. It derives the update host from the installed app (or `SWARMFILE_UPDATER_ENDPOINT`), downloads the **full installer** for the platform, and runs it (`.pkg`, `msiexec`, or `dpkg -i`), so it needs the same privileges a normal install does (`sudo /SYSTEM`).

On Windows, an update also refuses to run when the install directory is writable by non-administrators — the elevated helper stages and swaps files there, so a location any user can write is not a safe update target. If you hit that, reinstall to a protected location such as `Program Files`.

A repair **keeps your local state**: cached content, the queued-upload database, config, and sign-in token all live outside the application bundle (see [Uninstall & Clean Reinstall](#) for the exact paths). It is the supported way to move a `.dmg`-only macOS install onto the `.pkg` layout that installs the engine service and puts the CLI on `PATH`.

If the update manifest cannot be read, `--repair` reinstalls the **currently installed** version rather than guessing — so a repair never accidentally upgrades or downgrades a machine.

Where to go next

- [Supported Platforms & System Requirements](#) — what each installer contains.
- [Network Requirements](#) — the hosts above in firewall-allowlist form.
- [Deploying to Your Team](#) — mass deployment, silent install, and repair at scale.

Diagnostics & Error Reference

When something goes wrong, Swarmfile names the problem in one of a few places. This page maps the surfaces — doctor checks, status counters, health endpoints, and error codes — to plain meanings and next steps.

Start here

Surface	What it is	How to reach it
Desktop App → Settings → Diagnostics	Live checks, report file, Repair drive , Reinstall... , update check	The tray's settings
Status pill (header)	Mount location, sync state, the hub's reason for any pause	Click it
<code>swarmfile status</code>	The same snapshot as JSON or prose: counters, peers, bandwidth	CLI
<code>swarmfile-doctor</code>	Standalone probe suite that runs with no engine and no sign-in	CLI
<code>swarmfile doctor</code>	Runs the same probes through a live engine	CLI
Hub error codes	Machine-readable reasons on API refusals (bring-your-own scripts)	JSON bodies

`swarmfile-doctor` writes a JSON report on every run and returns a verdict:

```
swarmfile-doctor --json --report doctor-report.json # full detail
swarmfile-doctor --repair --yes                    # reinstall the package (keeps your queue)
```

Exit codes: **0** everything passed, **1** at least one failure, **2** warnings only. The report defaults to `~/.cache/swarmfile/last-doctor.json` (macOS/Linux) or `%USERPROFILE%\Swarmfile\last-doctor.json` (Windows).

Engine logs live in the cache directory (`~/.cache/swarmfile` on macOS and Linux, `%LOCALAPPDATA%\swarmfile` on Windows) as `engine.log`; the Desktop App writes `tray.log` beside it, and macOS mount-helper output lands in `~/Library/Logs/fuse-t/fuse-t.log`. Diagnostics has an **Open log directory** action. Increase detail with `RUST_LOG=swarmfile_engine=debug`.

Status counters

`swarmfile status` (and the Desktop App's panels) report these; the names are the JSON fields from `swarmfile status --json`:

Counter	Meaning	What to do
<code>pendingUploads</code>	Blocks queued on this machine awaiting cloud upload, plus closed files whose save is still running	Wait; check for connection trouble
<code>peerPendingUploads</code>	In-flight uploads reported by other machines on this branch	Usually informational
<code>pendingCreates</code>	Files/folders created locally and queued for the hub (an object with the count, stuck rows at six or more failed attempts, and the oldest age)	They land automatically; see Coding Agents
<code>createsReconciled</code>	Creates whose response was lost and whose registry row the engine then adopted by re-listing — the "dialog said the create failed but the project exists" case, recovered	Informational; non-zero is for support (<code>swarmfile status --json</code>)
<code>commitsPending</code>	Bytes uploaded but the commit hasn't landed (retrying, rate-limited, quota-held, or locked)	Check <code>quotaHeldCommits</code> and connectivity
<code>quotaHeldCommits</code>	Commits refused on billing grounds, retried on a 15-minute cadence; <code>spendCapHeldCommits</code> is the subset held by the org's spend cap, which the tray names ("Waiting on spend cap")	Raise the limit or free space; retries resume
<code>commitsRefused</code>	Commits the hub keeps refusing for other reasons (a subset of <code>commitsPending</code>)	Investigate the refusal; usually policy or conflict
<code>deferredOps</code>	Deferred rename/delete queue: ops applied locally and still owed to the hub (<code>queued</code>), how many are <code>held</code> (parked on a lock, a closed branch scope, or a child move still settling), the oldest op's age, and lifetime <code>confirmed</code> / <code>rolledBack</code> counts	Informational; <code>held</code> is the one to check — it makes no progress until the holder changes. See Working with Files
<code>abandonedGroups</code>	Upload groups that exhausted retries or hit a plan cap	<code>swarmfile sync stuck</code> lists them; re-save after fixing the cause
<code>failedWrites</code>	Guest-session writes permanently given up (guest replay only)	Re-open the guest session and retry
<code>conflicts</code>	Files blocked by an unresolved conflict. A blocked write reaches an application as a plain I/O error, so <code>status</code> counts it here instead of staying silent	Run <code>swarmfile conflicts</code> to list the paths and resolve
<code>activeLocks</code>	Write locks this engine currently holds	Informational
<code>checkoutLocks</code>	Long-lived per-file checkout locks (explicit, git-LFS) you hold	Release with <code>unlock</code> / <code>offline return</code> when done
<code>scopeReservations</code>	Offline scope reservations this engine holds — one row can cover a whole subtree; Pack & Go uses these rather than per-file locks	<code>swarmfile offline status</code>
<code>checkoutRenewFailures</code>	Renewal passes that re-issued nothing — reservations at risk of lapsing	Reconnect and re-prepare before the lease ends
<code>checkoutExpiringSoon</code>	Reservations expiring within 24 hours	Extend or return; see Working Offline
<code>peers</code>	LAN/office/seed peers currently reachable, with RTT	More peers = faster reads for shared files
<code>bandwidth.bytesIn</code> / <code>bandwidth.bytesOut</code>	Transfer volume for this session	Useful when a link is saturated
<code>blockStore</code>	Pinned blocks and local cache usage	<code>swarmfile cache status</code> / <code>hydrate status</code>
<code>auth</code>	Session state ("signed out" when there's no usable session)	Sign in again
<code>recentErrors</code>	A ring of the most recent engine errors	Often the fastest pointer to the root cause
<code>lastSequenceId</code>	Latest committed sequence known here	Diagnostic for history lag
<code>mounted</code> / <code>mountState</code> / <code>mountError</code>	Whether the mount is live and, if not, why (a stable code plus a user-facing sentence)	See the mount errors below

Health endpoints

There is **no default health port** — the server starts only when `SWARMFILE_HEALTH_PORT` is set (typical choices: `8080` in container deployments, `9100` for a seed node). It binds `0.0.0.0:<port>` and serves exactly:

- `GET /livez` — a cheap `{"status":"ok"}` for liveness probes (Kubernetes and friends). Allowed from anywhere.
- `GET /healthz` — the detailed snapshot (version, uptime, peer counts and RTT, pin/cache stats, encryption state). **Loopback-only:** a non-loopback caller gets `403 {"error":"loopback only"}`, because the detail exposes operational internals.
- Anything else — `404`.

Doctor checks

`swarmfile-doctor` reports one row per check. The whole report is capped at 10 seconds; if the cap fires, a synthetic `report` check fails, and a probe that couldn't complete counts as a failure rather than a silent pass. Each of the three network checks — `hub https /health`, `oidc token`, `cloud path` — retries one fast transient failure (a connection error or a 5xx) before it reports Fail, so a brief hub blip doesn't read as an outage; a definitive refusal like an expired token's `401` is reported on the first reply. The ones that name a common failure:

Check	Failing usually means
engine-running	No engine for this user; sign in / start the app, or expect the standalone checks below
install-layout	The engine isn't where the installer put it (moved out of the app bundle, or an older macOS install whose LaunchAgent still runs a stale copy); reinstall or use Repair . On macOS it also warns when <code>/usr/local/bin</code> is missing the CLI-tool links (a <code>.dmg</code> -only install the tray couldn't repair), since <code>swarmfile</code> is then not on <code>PATH</code> .
updater-helper-task	The elevated Windows updater helper is missing or mis-registered; reinstall (Windows)
dns <host>	DNS resolution for the hub/identity/storage host failed
hub https /health	The hub host is unreachable or blocked; proxy/firewall issue
hub deep health	The hub is up but one of its dependencies (its database, storage, the identity provider, or its own clock) is degraded
oidc token	The sign-in session expired, was revoked, or the account isn't a member of the configured org; sign in again
cloud path	Cloud storage credentials, presigning, or the block path failed for this org
iroh quic dial / mdns listen	Peer-to-peer path blocked (firewall, VLAN, VPN); falls back to the hub
host firewall / macos local network	OS firewall or macOS Local Network consent blocks LAN discovery
uploads in flight / in-flight streaming	Uploads or reads waiting on in-flight bytes exist — expected during activity, informative otherwise
fuse-t mount helper / mount	The macOS mount helper died, the mount point is wedged or absent, or a killed session left a stuck mount attempt behind — the engine clears those before each mount, so try Repair drive first; a stuck attempt that survives the repair needs a restart of the computer
duplicate mount stack	Two engines are competing for the same mount point
exclusive writes	Never fails — informational: names the drives in the opt-in exclusive-write mode
history retention	Never fails — informational: the active project's retention mode (every saved version kept and cold-archived, or named commits only). "Not checked" without a project scope
upload quota	Blocks are parked because the org is at a billing limit — see the count/bytes/age in the row
plan	Subscription inactive, a feature isn't in the plan (e.g. seed mode), or a metered dimension is $\geq 80\%$
pack and go policy	A <code>SWARMFILE_PACK_AND_GO_POLICY</code> value that isn't <code>allowed</code> / <code>disabled</code> — failing closed
previous run	The last engine run ended abnormally (panic or unexpected exit) — the row points at <code>engine.log</code>
url scheme	<code>swarmfile://</code> links (Open in app, sign-in, Explorer's Swarmfile menu) won't reach this install: no app is registered for the scheme, another app or install owns it, or (macOS) several copies of Swarmfile compete for it and a link may start an old one. Fix link handling re-registers this install; on Windows, a registration pointing at another install needs Repair from Settings > Apps

Engine and mount errors (what the app shows)

Code / message	Meaning	Fix
<code>not_signed_in</code>	The engine has no usable session	Sign in from the app
<code>not_a_member</code>	The session doesn't include the configured org/project	Switch organization or ask for an invite
<code>project_deleted</code>	The mounted project was permanently deleted	Mount another project; contact support about recovery within retention
<code>config_parse_error</code>	A <code>config.json</code> couldn't be read, so the engine came up without a project	The message names the file and the parse position — fix it and restart; see Engine Config File
<code>mount_failed / no_free_drive_letter</code>	Windows could not attach the mount	Free a drive letter or choose one in the app; <code>swarmfile mounts open auto</code>
<code>drive_letter_in_use</code>	The chosen letter (or Auto's candidates) are all taken	The message names what holds it; pick another
<code>mount_point_not_empty</code>	Windows refuses to mount over a non-empty folder	Choose an empty folder
<code>pack_and_go_disabled</code>	An administrator disabled bulk offline reservation	Use a per-file lock or ask your admin
<code>access_lease_expired</code>	The organization's offline access window elapsed — this computer hasn't reached Swarmfile within it, so local files are paused	Reconnect to the network; the drive resumes on its own (a restart won't help)
<code>confirmation_required</code>	A destructive CLI command was run without <code>--yes</code> in a non-interactive context	Re-run with <code>-y</code> after checking what it will delete
<code>usage_error</code>	A CLI invocation didn't parse (unknown command, flag, or argument); emitted as JSON when <code>--json</code> is present, with clap's <code>kind</code>	Fix the invocation — see Scripting against the CLI
<code>changelist_open</code>	A branch/project switch was refused because staged work is open	Submit, cancel, or <code>--park</code> first
<code>cross_project_mount_open</code>	An org switch was refused because an open drive points at a different project — that drive can't follow the engine to another org	Close that drive (the Desktop App names it in the switch confirm), then switch
<code>changelists_open</code>	<code>swarmfile mounts close</code> was refused because the mount has open changelists	Submit or cancel them, or pass <code>--force</code>
<code>is_default_mount</code>	You tried to close the boot mount on its own	Close another mount, or quit the engine instead
<code>shared_branch_context</code>	The mount shares the engine's branch context	Re-open it with <code>mounts open --branch <name></code> to make it independently switchable
<code>already_running / project_switch_running</code>	Another switch or long job of that kind is in flight	Wait for it to finish, then retry
<code>pending_create</code>	A lock/comment/share on a file that is still being created	Retry after it syncs (<code>swarmfile status</code> → <code>pendingCreates</code>)
<code>account_switch_requires_restart</code>	A runtime sign-in for a different account was refused	Restart the app (the tray offers the restart); the requested sign-in is kept and the restart completes the switch — the local cache identity is reconciled at boot, and the previous account's unsent work is quarantined, not lost
<code>account_switch_park_failed</code>	The requested account switch couldn't be saved for the restart (e.g. the disk is full)	Your current session is unaffected — free up disk space and try signing in again

Hub error codes (for scripts and integrations)

API refusals carry a machine-readable `code`. The ones most worth handling, in the order you'll meet them:

Code	Meaning	Typical handling
<code>bad_request</code> , <code>invalid_request</code> , <code>invalid_state</code> , <code>invalid_name</code> , <code>bad_cid</code> , <code>invalid_cid</code> , <code>bad_hash</code>	Malformed or out-of-order request	Fix the request; don't retry blindly
<code>credential_invalid</code> , <code>email_verification_required</code> , <code>email_not_verified</code> , <code>session_expired</code> , <code>session_revoked</code> , <code>unauthenticated</code> , <code>terms_acceptance_required</code>	Auth/session problem	Sign in again; verify email for share actions; accept updated terms
<code>api_key_*</code> (<code>api_key_grant</code> , <code>api_key_org_mismatch</code> , <code>api_key_project_mismatch</code> , <code>api_key_scope_forbidden</code>)	API-key scope doesn't cover this operation	Use a key scoped to the project/org, or mint a new one
<code>forbidden</code> , <code>not_member</code> , <code>admin_role_required</code> , <code>idp_binding_rejected</code>	Not permitted for this account/role	Escalate to an admin
<code>root_create_restricted</code>	Creating or moving an entry directly at the root of a protected project, by a caller who isn't the project creator and holds no project-wide <code>write</code> grant (a folder-scoped grant, even <code>admin</code> , doesn't qualify)	Create inside an existing folder, or ask an owner or the project creator; a held <code>create-queue</code> save resumes once access changes, and a <code>git push</code> that would add a root name is refused the same way — its error names the remedy (move the files inside a folder, or gain project-wide write)
<code>org_mfa_required</code> , <code>org_email_verification_required</code>	The org's authentication policy refuses this credential	Enroll an authenticator app / verify your email under Account settings, then sign in again — the refusal body's <code>enrollUrl</code> points there. Don't retry unchanged; a personal access token must be re-minted from a compliant session
<code>owner_mfa_required</code> , <code>mfa_status_unavailable</code>	Turning on "require MFA" needs the acting owner's own confirmed factor (or the factor-status probe was unreachable)	Enroll an authenticator app first, then retry the policy change
<code>identity_conflict</code>	An external-IdP sign-in conflicts with an existing Swarmfile account: its subject names a first-party user, or a principal provisioned by a different issuer	Contact the organization's administrator; sign in through the org's own SSO rather than a built-in Swarmfile account
<code>entry_not_found</code> , <code>entry_deleted</code> , <code>project_not_found</code> , <code>project_deleted</code> , <code>path_not_found</code> , <code>not_a_file</code> , <code>source_gone</code> , <code>target_gone</code>	The target is gone	Re-list; treat as terminal
<code>exists</code> , <code>name_conflict</code> , <code>id_conflict</code> , <code>name_taken</code>	Collision or illegal name	Rename / resolve the conflict
<code>entry_uploading</code>	The entry has no committed version yet	Wait for upload, then retry
<code>unscoped_project_route</code>	The client called an org URL that needs <code>/projects/:id</code> — an out-of-date desktop app does this	Update Swarmfile; the desktop app holds the affected save and reports it once instead of retrying
<code>stream_retired</code>	The notification/activity stream isn't supported; clients poll instead	Poll <code>GET /orgs/:orgId/notifications</code> (and <code>/activity</code>); don't reconnect the stream
<code>stage_discarded</code>	A <code>POST /jobs/:id/retry</code> was refused because the job's staged operation was already discarded	Start the operation again — there is nothing left to retry
<code>not_cancellable</code>	<code>POST /jobs/:id/cancel</code> (and <code>swarmfile op cancel</code>) refused: this operation can't stop part-way — a commit, merge or delete completes or fails on its own	Let it finish; follow it with <code>op wait</code>
<code>too_late</code>	<code>POST /jobs/:id/cancel</code> refused: the staged revert/checkout has already begun applying (including the moment after it published), so there is no unpublished stage to discard	Nothing to stop — the commit lands; the operation ends with its real outcome (<code>op status</code> / <code>op wait</code>), not <code>canceled</code>
<code>chunk_missing</code> , <code>cid_mismatch</code> , <code>checksum_mismatch</code> , <code>size_mismatch</code> , <code>manifest_unreadable</code> , <code>manifest_too_large</code> , <code>tree_object_corrupt</code> , <code>tree_object_missing</code> , <code>commit_object_corrupt</code> , <code>commit_object_missing</code>	Content integrity/structure problem	Re-upload the file; report if persistent
<code>conflict_detected</code> , <code>conflict_unresolved</code> , <code>merge_conflict</code> , <code>stale_merge</code>	Concurrent edits, or the target moved since the merge/conflict was prepared	Resolve the conflict (<code>swarmfile conflicts resolve</code>), then retry
<i>(no code; the envelope's error text says <code>target_changed</code> / <code>merge_candidate_changed</code>)</i>	The file changed again after the conflict was recorded, or the automatic merge result changed after you previewed it	Re-read the file and resolve against the newer version; run the merge attempt again before accepting. Match the text, not a <code>code</code> — the hub sends these two without one

Code	Meaning	Typical handling
locked	Someone else holds the file, or a machine's Pack & Go scope reservation covers it	Wait, or request an unlock (<code>swarmfile unlock-request</code>); a scope reservation is released by its owner with <code>swarmfile offline return</code>
commit_mode_locked , policy_enforced	An admin lock requires a different commit mode, or refuses the setting	Check the enforced mode with <code>swarmfile config get-commit-mode</code> , then use it
commit_race , head_changed , merge_in_progress , fork_in_progress	History moved under you, or a merge/fork is running	Re-read, wait, then retry; when a project job holds the same operation, the refusal names that job — retry or cancel it first
branch_not_found , branch_mismatch , branch_locked , branch_has_open_mr , branch_has_active_children , merge_main	Branch rules prevented the action	Use the allowed path (merge the MR, delete children first, never merge into main directly)
not_approved , self_review , is_draft , not_draft , protected_branch	Merge-request gates	Get approvals, un-draft, don't self-review
project_archived , storage_exceeded , quota_exceeded , spend_cap_reached , plan_restricted , keys_cap , subscription_inactive , too_many_entries , too_many_descendants , registry_too_large , batch_too_large , body_too_large , merge_too_large , too_many_cids	Plan or size limit	Free space, upgrade, or split the request/project; <code>merge_too_large</code> means the merge exceeded the single-project merge capacity — land the branch in smaller pieces — otherwise retry after the allowance changes. <code>spend_cap_reached</code> is the dollar ceiling: the envelope carries <code>capCents / usedCents / remainingCents</code> (and <code>keyLimitCents / keyId</code> when a per-key budget was the binding wall) — raise the cap or free space and held writes resume on their own
project_storage_full	This project hits its metadata ceiling (files and versions combined)	Delete or purge files, delete stale branches, shorten history retention, or switch the project to Named commits only (the same guidance the owner notification carries) — see Project storage limit
checkout_required	The org hasn't completed subscription checkout	Finish checkout in Billing before creating projects
rate_limited	Request budget exceeded	Back off and retry; see Rate limits
grace_period_inaccessible	The org is in the post-unpaid 28-day window	Resubscribe — see Data Portability
quarantined	Ransomware detection quarantined the account; writes stop until an owner releases it (on Windows the drive stays mounted meanwhile, with saves held locally)	Review and release from the quarantine panel — see Ransomware quarantine
project_key_unavailable , project_key_rotation_unsupported	Encryption key unavailable, or the project can't rotate (plaintext storage, or it has reached the maximum key generations)	Check the project's encryption tier; contact support if the cap is reached
stale_generation , incomplete_recipients , unexpected_recipients , key_rotation_cap_reached	E2E key rotation conflicts: another device rotated first; the caller's wrap set was missing (or named non-) current recipients; or the 64-generation cap was reached	Re-run the rotation (the engine re-fetches the recipient set and retries once automatically); an owner should confirm the device roster before rotating
e2e_upload_unsupported , e2e_unsupported , e2e_not_supported , encryption_tier_not_public	The operation isn't allowed for this tier	Use the supported tier (e.g. LFS is rejected on E2E)
jurisdiction_incompatible	The destination isn't in the pinned jurisdiction — the org's, or a project's own	Use storage/a destination in the pinned region
residency_unavailable , residency_storage_unavailable , jurisdiction_unverifiable , residency_unverifiable	A per-project residency pin was refused: this deployment can't route one, the region's storage isn't ready yet, or the check couldn't be completed just now	Retry shortly (a region still provisioning becomes ready on its own); if it persists, contact support
share_too_large , block_not_in_share , block_session_expired , raw_byte_budget_exceeded	Share-link limits or a public-streaming budget	Reduce scope, create a new share, or wait out the budget
too_many_webhooks , blocked_target , invalid_url	Webhook limits or an unsafe target URL	Delete unused webhooks; use a public HTTPS endpoint
preview_budget_exhausted , preview_too_large , preview_type_unsupported	Preview generation budget, size, or type limit	Wait for the daily reset; serve the file directly or download it
too_many_active_jobs , too_many_concurrent_reads	Temporary concurrency limit	Back off and retry
ci_not_available , ci_disabled , ci_branch_unprotected , ci_e2e_unsupported	Hosted CI is refused for this org, project or branch: a project owner hasn't enabled it, the branch isn't protected and unprotected runs aren't opted in, or the project is end-to-end encrypted (the runner holds no E2E keys)	Enable it as an owner (<code>swarmfile ci enable</code>); protect the branch, or opt in with <code>—allow-unprotected</code> ; E2E projects can't run hosted CI — see Hosted CI

Code	Meaning	Typical handling
<code>ci_workflow_invalid</code> , <code>ci_yaml_invalid</code> , <code>ci_expression_invalid</code> , <code>ci_jobs_too_many</code> , <code>ci_secret_undeclared</code>	<code>.swarmfile/ci.yml</code> didn't compile (the error carries the line), the expansion passed 256 jobs, or a job referenced a secret it didn't declare	Fix the workflow; declare every secret a job reads in its <code>secrets:</code> list — see Hosted CI
<code>invite_unavailable</code> , <code>org_not_empty</code> , <code>installation_exists</code> , <code>already_provisioning</code> , <code>storage_provisioning</code>	Admin lifecycle state	Complete or wait for the pending operation
<code>announce_unverified</code>	A seed node's announce failed node authentication	Re-enroll the node; see Self-Hosted Seed Nodes
<code>org_migration_frozen</code> , <code>project_unreachable</code> , <code>async_unavailable</code>	A migration/maintenance state	Retry after it completes
<code>not_found</code> , <code>forbidden</code>	Generic — the meaning depends on the operation	Use your caller's contextual fallback; don't guess

The full set is generated from the hub's own routes and shipped in the hub's published API schema and OpenAPI contract; an unrecognized code is preserved verbatim rather than dropped, so a script can always log it.

Exit codes

Binary	Codes
<code>swarmfile</code>	<code>0</code> success; <code>1</code> failure (unreachable engine, transport error, no product equivalent); <code>2</code> actionable refusal or a usage error — under <code>--json</code> the two are distinguishable by <code>code</code> (<code>usage_error</code> vs the refusal's code); <code>3</code> a long operation is still running (returned by <code>--no-wait</code> , or after Ctrl-C detaches the command) — follow it with <code>swarmfile op status / op wait</code> ; see Scripting against the CLI
<code>swarmfile-doctor</code>	<code>0</code> all passed; <code>1</code> ≥1 failure; <code>2</code> warnings only. A usage error under <code>--json</code> prints the shared <code>{"code":"usage_error",...}</code> body, as every standalone tool now does
<code>swarmfile-runner</code>	The engine's code (the alias re-execs it); <code>127</code> if the sibling engine is missing
<code>swarmfile-migrate</code> , <code>swarmfile-search</code> , <code>swarmfile-seed</code>	<code>0</code> success; <code>1</code> failure (a failed verification or upload, an unreachable hub); <code>2</code> a malformed invocation — the shared <code>{"code":"usage_error",...}</code> body under <code>--json</code>
<code>swarmfile-verify-history</code>	<code>0</code> the loaded history verified; <code>1</code> verification found tampering or a gap; <code>2</code> the check couldn't run (bad arguments, unreachable hub, auth error) — see the CLI reference
<code>swarmfile-lfs-transfer</code>	Not a scripted CLI: git-LFS spawns it for one operation and reads its protocol replies. It exits <code>1</code> only when it can't start

Where to go next

- [Troubleshooting](#) — symptom-first workflows.
- [Network Requirements](#) — the hosts and ports the checks probe.
- [Environment Variable Index](#) — `SWARMFILE_HEALTH_PORT`, `RUST_LOG`, and the diagnostic switches.

Environment Variable Index

Almost every engine setting has an environment-variable form, and some settings exist **only** as environment variables. This page is the index: what each `SWARMFILE_*` variable does, its default, and whether the same setting can be put in `config.json` instead.

Use environment variables when a launcher can set them (systemd, Docker, a scheduled task, a CI job). Use `config.json` when it can't — notably macOS launchd agents and Windows services, which don't inherit a shell environment. See [Engine Config File](#) for the file's precedence rules and the list of keys a packaged install manages itself.

To see what a running install actually resolved — secrets redacted — run `swarmfile-engine --print-env-vars`.

Conventions

- Booleans accept `true / 1` to enable; anything else (including an empty string) is false. A few "disable with" flags are the reverse: only `false / 0` disables them.
- Empty strings are treated as unset for credentials, ids, paths, and tool commands.
- Values set here beat the `config.json` file; both beat built-in defaults.
- Variables marked **advanced** are safe but rarely needed — reach for them while diagnosing, not as standard configuration.
- Variables marked **diagnostic only** are for support sessions and are noisy.

Credentials, identity, and routing

Variable	Default	Purpose	File key
<code>SWARMFILE_HUB_URL</code>	<code>https://hub.swarmfile.com</code> on a packaged install	Hub API base URL.	<code>hub_url</code>
<code>SWARMFILE_SITE_URL</code>	derived from the hub host (<code>hub.X</code> → <code>X</code>) for engine-printed links; the compiled build value for the tray	Web-app base URL for links: the tray's in-app and dashboard links, and the links the engine's git-push helper prints. Wins over the installer's <code>site_url</code> . With neither, the git-push helper derives it from the hub host (and omits the link when even that can't be derived); the tray uses the value compiled into the build.	<code>site_url</code>
<code>SWARMFILE_ORG_ID</code>	<i>(unset)</i>	Org whose hub routes are called.	<code>org_id</code>
<code>SWARMFILE_PROJECT_ID</code>	<i>(unset)</i>	Project every metadata operation is scoped to (required for headless nodes).	<code>project_id</code>
<code>SWARMFILE_BRANCH</code>	the project's default branch (each project names it at creation — commonly <code>main</code>)	Branch this mount starts on.	<code>branch</code>
<code>SWARMFILE_OFFICE_ID</code>	<code>default</code>	Peer-grouping label for same-office LAN behavior.	<code>office_id</code>
<code>SWARMFILE_OIDC_ISSUER</code>	<code>https://id.swarmfile.com</code> on a packaged install	Identity provider the engine signs in against.	<code>oidc_issuer</code>
<code>SWARMFILE_OIDC_CLIENT_ID</code>	<code>swarmfile-desktop</code> on a packaged install	OIDC client id.	<code>oidc_client_id</code>
<code>SWARMFILE_OIDC_REFRESH_TOKEN</code>	<i>(unset)</i>	Non-interactive sign-in for headless engines; the rotated token is persisted under the cache dir.	—
<code>SWARMFILE_API_KEY</code>	<i>(unset)</i>	Project-scoped <code>sf_key_...</code> credential; wins over a refresh token if both are set.	—
<code>SWARMFILE_NO_KEYCHAIN</code>	<i>(unset)</i>	Use a file-backed secret store instead of the OS keystore — required on headless or agent-launched engines that can't show a Keychain prompt. Advanced	—
<code>SWARMFILE_SWARM_KEY</code>	<i>(unset)</i>	Extra deployment-wide entropy for the private peer-to-peer swarm.	—
<code>SWARMFILE_USER_ID</code>	<code>\$USER / \$USERNAME</code> , else <code>anonymous</code>	Identity recorded for locks and presence.	—
<code>SWARMFILE_MACHINE_ID</code>	OS hostname (tray: hash of hostname + data dir)	Machine identity for locks and presence.	—
<code>SWARMFILE_NODE_AUTH</code>	<code>false</code>	For orgs that enforce P2P node authentication: refuse all peers until the hub confirms the approved roster.	<code>node_auth</code>
<code>SWARMFILE_LAN_FROM_OFFICE</code>	<code>false</code>	Treat same-office peers discovered through the hub as LAN peers when mDNS is blocked.	<code>lan_from_office</code>

Paths and local data

Variable	Default	Purpose	File key
SWARMFILE_CONFIG	<i>(unset)</i>	Explicit single <code>config.json</code> path; bypasses the install + per-user merge.	—
SWARMFILE_DATA_NAMESPACE	<i>(empty)</i>	Suffixes the data dir, mount, socket, and lock so a second install can coexist.	<code>data_namespace</code>
SWARMFILE_CACHE_DIR	<code>~/ .cache/swarmfile</code> (macOS/Linux), <code>%LOCALAPPDATA%\swarmfile</code> (Windows)	Block/metadata cache, auth token, control socket, logs.	—
SWARMFILE_CACHE_MAX_BYTES	10737418240 (10 GiB)	Local block-cache cap.	<code>cache_max_bytes</code>
SWARMFILE_MOUNT_POINT	<code>~/Swarmfile</code> (Unix), <code>S:</code> (Windows)	Where the boot mount attaches.	<code>mount_point</code>
SWARMFILE_IROH_DIR	<code><cache_dir>/iroh</code>	Peer-to-peer node identity and blob store.	—
SWARMFILE_READ_ONLY	<code>false</code>	Force a kernel read-only mount (guest sessions).	<code>read_only</code>
SWARMFILE_SESSION_KIND	<code>owner</code>	Session role (<code>owner</code> / <code>member</code> / <code>guest</code>).	<code>session_kind</code>
SWARMFILE_HUB_ONLY	<code>false</code>	Cloud-only mode: no QUIC, no mDNS, no peer traffic. An org policy can also force it on.	<code>hub_only</code>

Networking and peer discovery

Variable	Default	Purpose	File key
SWARMFILE_IROH_BIND_ADDR	OS-assigned ephemeral port	Fixed QUIC bind address, for a firewall that needs a predictable port.	<code>iroh_bind_addr</code>
SWARMFILE_IROH_PUBLIC_ADDR	<i>(unset)</i>	Public address advertised to the hub when behind NAT.	—
SWARMFILE_IROH_PEER_BUDGET_SECS	30	Aggregate time budget for peer fetches on interactive paths. Advanced	—
SWARMFILE_IROH_PEER_BUDGET_BACKGROUND_SECS	10	Same budget for background prefetch. Advanced	—
SWARMFILE_HEALTH_PORT	<i>(unset — health server off)</i>	Port for the <code>/healthz</code> and <code>/livez</code> endpoints.	—
SWARMFILE_MDNS_INTERFACES	auto-selected LAN interfaces	Comma-separated allowlist mDNS must use exclusively.	—
SWARMFILE_MDNS_EXCLUDE_INTERFACES	<i>(unset)</i>	Comma-separated interfaces mDNS must never use.	—
SWARMFILE_MDNS_SERVICE_TYPE	<code>_swarmfile._udp.local.</code>	Override the mDNS service type so an engine stays off the shared type (for isolated test networks). Must be a valid <code>_<name>._udp.local.</code> type; an invalid value is logged and disables mDNS (the engine continues with hub-only peer discovery) rather than silently using the default type. Advanced	—
SWARMFILE_WORKSHARING_FAIL_CLOSED	<code>false</code>	Refuse an exclusive-open check when the hub is unreachable, instead of allowing the open.	—
SWARMFILE_GOSSIP_ATTESTATION_MODE	<code>off</code> (<code>off</code> / <code>log</code> / <code>enforce</code>)	Hub-attested gossip enforcement. <code>log</code> records what enforcement would refuse; <code>enforce</code> refuses unverified gossip. Off by default. Advanced	—

Encryption and erasure coding

Variable	Default	Purpose	File key
SWARMFILE_ENCRYPTION_ENABLED	<code>true</code>	Master switch for block encryption.	—
SWARMFILE_ENCRYPTION_KEY	<i>(unset)</i>	Hex 32-byte static master key (no hub key fetch).	—
SWARMFILE_ENCRYPTION_PROJECT_ID	falls back to <code>SWARMFILE_PROJECT_ID</code>	Project whose hub-managed key is fetched.	<code>encryption_project_id</code>
SWARMFILE_EC_ENABLED	<code>auto</code>	Erasure coding: <code>on</code> / <code>true</code> , <code>off</code> / <code>false</code> , or <code>auto</code> (adapt per save).	—
SWARMFILE_EC_LAN_PLACEMENT	<code>false</code>	Deliberately spread erasure shards across LAN peers.	<code>ec_lan_placement</code>
SWARMFILE_EC_LAN_PLACEMENT_REPLICATION	3	Replication factor K per shard. Advanced	—
SWARMFILE_EC_LAN_PLACEMENT_POLL_INTERVAL	300 (seconds)	Placement reconcile interval. Advanced	—

Seed mode and CI runner

Variable	Default	Purpose	File key
SWARMFILE_SEED_MODE	false	Run as a NAS/seed node: fetch every block into a warm cache, mirror to cloud storage, no mount. A headless seed also needs SWARMFILE_ORG_ID / _PROJECT_ID and a credential; the dashboard's Settings → Office caches wizard generates the full file.	seed_mode
SWARMFILE_SEED_POLL_INTERVAL	60 (seconds)	How often a seed polls the hub for metadata changes.	—
SWARMFILE_SEED_UPLOAD_CONCURRENCY	8	Concurrent cloud uploads from a seed.	—
SWARMFILE_RUNNER_MODE	set by swarmfile-runner	Headless CI runner; watches SWARMFILE_BRANCH for commits/tags matching .swarmfile/runner.yml . Refuses to execute the rule file on a branch without an effective branch protection unless SWARMFILE_RUNNER_ALLOW_UNPROTECTED=1 is set (trusted/dev runners only).	—
SWARMFILE_RUNNER_ALLOW_UNPROTECTED	(unset — protected branches only)	Let a headless runner execute its rule file on a branch with no effective branch protection. Trusted/dev runners only; leave unset in production.	—

Create queue, streaming, and mount I/O

Variable	Default	Purpose	File key
<code>SWARMFILE_PENDING_CREATE_WAIT_SECS</code>	8 (0–120; 0 refuses immediately)	How long an operation on a still-syncing file waits before answering <code>409 pending_create</code> .	—
<code>SWARMFILE_CREATE_DISPATCH_CONCURRENCY</code>	8 (1–64)	Creates sent to the hub at once. Advanced	—
<code>SWARMFILE_CREATE_COALESCE</code>	on (0 / false / off disables)	Attach the first save to a create instead of sending the create empty. Advanced feature switch	—
<code>SWARMFILE_CREATE_HOLD_SECS</code>	2 (0–10; 0 disables holding)	Max wait for the first save of an open-but-unsaved file. Advanced	—
<code>SWARMFILE_CREATE_HOLD_UPLOADING_SECS</code>	60 (0–600)	Same, when that first save is already uploading. Advanced	—
<code>SWARMFILE_CREATE_BATCH</code>	on (0 / false / off disables)	Batch creates through the hub's batch route. Advanced feature switch	—
<code>SWARMFILE_ARRIVING_REPORTS</code>	on (0 / false / off disables)	Publish the "N files arriving" presence hint while creates are in flight. Advanced feature switch	—
<code>SWARMFILE_STREAM_BUFFER_SECS</code>	30 (seconds)	Read-ahead a <code>fetch</code> <code>—follow</code> job keeps buffered. Advanced	—
<code>SWARMFILE_STREAM_READ_WAIT_SECS</code>	15 (seconds)	How long one mount read waits for in-flight bytes before a retryable error. Advanced	—
<code>SWARMFILE_STREAM_WAIT_CONCURRENCY</code>	4 (1–64)	How many reads may be waiting at once. Advanced	—
<code>SWARMFILE_READAHEAD_CONCURRENCY</code>	12 (1–64)	Concurrent mount read-ahead fetches for an established sequential run; a run's first, speculative batch stays at 4. Advanced	—
<code>SWARMFILE_READAHEAD_MAX_MB</code>	(per file type)	Mount read-ahead window ceiling, in MiB; 0 disables read-ahead entirely. Advanced	—
<code>SWARMFILE_READAHEAD_BUFFER_SECS</code>	(per file type; 15 s for video containers)	Seconds of playback the rate-adaptive read-ahead ceiling keeps ahead; 0 = the flat profile ceiling. Overrides the file value. Advanced	<code>readahead_buffer_secs</code>
<code>SWARMFILE_MOUNT_ATTEMPT_TIMEOUT_SECS</code>	120 (0 disables the bound)	Wall-clock bound on a mount attempt so a wedged driver can't hang the app. Advanced	—
<code>SWARMFILE_FS_CALLBACK_RESCUE_SECS</code>	(unset — off)	Windows only. When a drive operation has been stuck this many seconds (30–3600), reset that drive the way Repair drive does, so the app waiting on it is released; at most once per drive every 10 minutes. Unsaved changes in the app that hung may be lost. Off by default: a stuck operation is only logged. Advanced	—
<code>SWARMFILE_FS_CALLBACK_STUCK_SECS</code>	30 (5–3600)	Windows only. A callback running longer than this is logged as stuck (the WARN the rescue knob reacts to). Advanced	—
<code>SWARMFILE_FS_CALLBACK_WATCHDOG_TICK_SECS</code>	10 (5–3600)	Windows only. How often the callback watchdog scans for stuck operations. Advanced	—
<code>SWARMFILE_MOUNT_OPEN_WAIT_SECS</code>	20 (0 = don't wait)	How long <code>mounts open</code> waits for the platform mount to settle. Advanced	—
<code>SWARMFILE_PACK_AND_GO_POLICY</code>	allowed ; any unrecognized value disables Pack & Go (and is logged at startup, plus a <code>swarmfile-doctor</code> warning)	Operator gate for bulk offline reservation — see Working Offline .	—
<code>SWARMFILE_GIT_ACCEPT_PROVISIONAL</code>	false	The <code>swarmfile://</code> remote helper accepts provisional (derived-but-unconfirmed) commit ids when resolving a shallow clone's boundary — needed for <code>—depth</code> on a project whose commits only one person derives. Unset, a boundary id must be confirmed. Advanced	—
<code>SWARMFILE_DEFERRED_RENAME</code> (alias <code>SWARMFILE_DEFERRED_NAMESPACE</code>)	(unset — on)	Apply non-displacing/displacing renames and file/empty-directory deletes of confirmed entries locally and queue them, while hub confirmation, retraction and refresh run in the background. Set to 0 / false / no / off to opt out; any other value is ignored (and logged) so the config file/default applies. Also settable as <code>deferred_namespace</code> in <code>config.json</code> (environment beats file; if both spellings are set, <code>SWARMFILE_DEFERRED_RENAME</code> wins). Advanced	<code>deferred_namespace</code>

Bandwidth, QoS, and sync behavior

Variable	Default	Purpose	File key
SWARMFILE_THROTTLE_ENABLED	false	Enable bandwidth throttling; set a site bandwidth first.	throttle_enabled
SWARMFILE_SITE_DOWNLOAD_BANDWIDTH_MBPS	(unset)	Site download capacity.	site_download_bandwidth_mbps
SWARMFILE_SITE_UPLOAD_BANDWIDTH_MBPS	(unset)	Site upload capacity.	site_upload_bandwidth_mbps
SWARMFILE_OFFICE_HOURS_START	08:00	Start of the office-hours window (local HH:MM).	office_hours_start
SWARMFILE_OFFICE_HOURS_END	18:00	End of the office-hours window.	office_hours_end
SWARMFILE_OFFICE_HOURS_WAN_PCT	20	WAN percentage allowed during office hours.	office_hours_wan_pct
SWARMFILE_OFF_HOURS_WAN_PCT	80	WAN percentage allowed outside them.	off_hours_wan_pct
SWARMFILE_DSCP_INTERACTIVE	34	DSCP marking for interactive traffic. Advanced	—
SWARMFILE_DSCP_BACKGROUND	0	DSCP marking for bulk traffic. Advanced	—
SWARMFILE_XREF_PREFETCH	true	Prefetch .dwg / .max external references on open.	—
SWARMFILE_SYNC_IGNORE_ENABLED	true	Honor .gitignore / .swarmfileignore exclusion.	sync_ignore_enabled
SWARMFILE_CHANGESET_IDLE_SECS	300 (0 disables session minting)	Idle timeout before a staged changeset commits.	changeset_idle_secs
SWARMFILE_COMMIT_LINGER_MS	0 (off; clamped to 2000)	Milliseconds to gather a burst of ready commits into one batch before committing. Advanced	—
SWARMFILE_RWS_TTL_DAYS	7	How long a read or written entry stays in the recent-working-set list that offline writes are admitted against. Advanced	—
SWARMFILE_AUTO_OPEN_CHANGELIST_FOR_PREFERENCE	false	Open a changelist on boot when the user's commit-mode preference asks for it.	auto_open_changelist_for_preference
SWARMFILE_ATTR_CACHE_SECS	2 (0 keeps FUSE-T's own default)	macOS/Linux attribute-cache ceiling.	attr_cache_secs
SWARMFILE_MERGETOOL_CMD	(unset — git config fallback)	External merge-tool command template.	mergetool_cmd
SWARMFILE_MERGETOOL_NAME	(unset)	Display name for that tool.	mergetool_name

Diagnostics and telemetry

Variable	Default	Purpose	File key
SWARMFILE_DOCTOR_PERIOD_SECS	300 (0 disables)	Period of the engine's background connectivity self-check.	—
SWARMFILE_MEMBERSHIP_WATCHDOG_SECS	5	How often the engine's membership-enforcement loop runs locally (flag checks, purge/close retries). A definitive refusal from any hub call sets the revocation immediately, so this interval affects how long an unobserved loss can linger, not detection itself; for an unmounted project, detection waits for the next hub contact (the opt-in background probe, SWARMFILE_BACKGROUND_REVOCATION_PROBE , runs at its own 300 s cadence). 0 /unset falls back to 5 .	—
SWARMFILE_OTEL_ENDPOINT	(unset — local collector in a diagnostic build)	OTLP collector endpoint for engine traces. Operator builds only — shipped desktop installs omit trace export entirely.	—
SWARMFILE_OTEL_SERVICE_NAME	swarmfile-engine	Service name reported to that collector (same operator-build caveat).	—
SWARMFILE_UPLOAD_TRACE	off (1 / true on)	Promote per-attempt upload logging to warn , including the hub's refusal body. Diagnostic only	—
SWARMFILE_FUSE_DEBUG	off	libfuse protocol tracing (macOS). Diagnostic only	—
SWARMFILE_WINFSP_DEBUG	off	WinFsp protocol tracing (Windows). Diagnostic only	—
SWARMFILE_UPDATER_ENDPOINT	installed config	Update manifest endpoint that swarmfile-doctor --repair repairs from.	tray config

Installers and the git-LFS agent

These are consumer/operator knobs rather than engine settings.

Variable	Default	Purpose
SWARMFILE_FW_PROFILES	3 (Domain + Private; bitmask: 1 Domain, 2 Private, 4 Public)	Windows MSI property controlling which firewall profiles get the LAN rules, e.g. <code>msiexec /i Swarmfile-x64.msi SWARMFILE_FW_PROFILES=7 /qn</code> .
SWARMFILE_LFS_VERSION	latest release	Pin the git-LFS agent version installed by the agent installer.
SWARMFILE_LFS_BINDIR	~/.local/bin (Unix), %LOCALAPPDATA%\Swarmfile\bin (Windows)	Where the LFS agent installer puts the binary.
SWARMFILE_LFS_REPO	swarmfile/swarmfile-lfs	Release repo the installer pulls from.
SWARMFILE_LFS_DOWNLOAD_BASE	provider default	Mirror/CDN base for agent downloads.
SWARMFILE_LFS_API_URL	provider default	Release-listing API override for a mirror.
SWARMFILE_LFS_ALLOW_UNVERIFIED	unset (1 = allow)	Proceed when a mirror strips .sha256 checksums. Security-relevant — only for a mirror you trust.

Standard environment variables the product honors

Variable	Effect
RUST_LOG	Log filter for the engine, CLI, and Desktop App backend (for example <code>RUST_LOG=swarmfile_engine=debug</code>).
HOME / XDG_CONFIG_HOME	Base for the cache dir and the per-user <code>config.json</code> on macOS/Linux.
LOCALAPPDATA / USERPROFILE	Base for the data dir and config discovery on Windows.
USER / USERNAME	Fallback identity when <code>SWARMFILE_USER_ID</code> is unset.
NO_COLOR	Disable ANSI color in <code>swarmfile-doctor</code> output.

Where to go next

- [Engine Config File](#) — the file form of these settings, precedence, and live vs restart-required keys.
- [Network Requirements](#) — what the networking variables above connect to.
- [Diagnostics & Error Reference](#) — the doctor checks that read these values.

Engine Config File (`config.json`)

The Swarmfile engine reads its startup configuration from these sources, in this order of precedence:

1. **Environment variables** (`SWARMFILE_*`) — always win, with one exception: `data_namespace` on a packaged install, where the bundled install config wins even over the env var (see "A note on environment-selection keys" below).
2. **A stored OIDC session** (*applies to `oidc_issuer` / `oidc_client_id` only*) — after an interactive sign-in, the engine remembers the issuer and client id it signed in against; those beat a hand-edit of `oidc_issuer` / `oidc_client_id` in the file. Editing either key on a signed-in install can appear to do nothing until you sign in again.
3. **The `config.json` file** — the subset of settings documented below.
4. **Built-in defaults** — a released build ships pointing at Swarmfile's hosted service, so an unconfigured engine connects to production rather than nothing.

Most people never touch this file: the Desktop App writes the keys it manages for you, and a normal sign-in fills in the rest. You edit `config.json` directly for **headless and unattended setups** — a NAS/seed node, a render-farm worker, a VM, or any machine that boots the engine from autostart rather than a signed-in Desktop App session — where there's no UI to click and, on macOS/Windows, no shell environment for the engine to inherit exported variables from.

Where the file lives

Platform	Default path
macOS / Linux	<code>~/config/swarmfile/config.json</code> (respects <code>\$XDG_CONFIG_HOME</code>)
Windows	<code>%LOCALAPPDATA%\swarmfile\config.json</code>

That per-user path is only half the picture. A packaged install also ships its own **install config** — read-only, written by the installer, and living next to the binary (on macOS, inside the bundle's `Contents/Resources/config.json` ; on Linux's `.deb` , `/etc/swarmfile/config.json`). At startup the engine merges the two, generally with the per-user file winning per key — except for a small, fixed set of keys a packaged install manages itself, where the install config wins instead. See "A note on environment-selection keys" below for exactly which keys, and why, before hand-editing any of them.

To point the engine at a **specific file** instead of this two-file merge — e.g. one config per project on a multi-tenant node — set `SWARMFILE_CONFIG=/path/to/config.json` . That bypasses the merge entirely and reads only that one file.

Editing safely

A few things are worth knowing before you hand-edit this file:

- **Restart the engine after editing.** The file is read once at startup. A handful of values can also be applied to a running engine through the control socket with no restart: `cache_max_bytes` , `readahead_buffer_secs` , `throttle_enabled` and the four office-hours keys, `office_id` , `lan_from_office` , and `branch` . Everything else is read only at startup (including `seed_mode` , `hub_only` , `ec_lan_placement` , `changeset_idle_secs` , `attr_cache_secs` , `mount_point` , `read_only` , and `session_kind`). `sync_ignore_enabled` updates live and then drains and restarts the engine. Either way, a change to the **file itself** needs a restart to be read.
- **A single syntax error discards the *whole* file — and says so.** If the JSON doesn't parse, the engine logs a warning, ignores that file, and keeps going with whatever the other sources provide. On a packaged install that means the bundled install config and environment still apply — but a broken per-user file still drops your `org_id` and `project_id` , so the drive has nothing behind it. Instead of a generic "no project" prompt, the Desktop App's status then names the file and the parser's position (for example `config.json: expected value at line 1 column 3`), because a project picker can't help when the org id lived in the file that failed to parse. (With `SWARMFILE_CONFIG` , there is no other file, so a bad file leaves the engine on built-in defaults.) A JSON syntax error — a trailing comma, say — silently drops those along with everything else in the file. Validate your JSON before restarting.
- **Unknown keys are named, not ignored in silence.** A valid file with a misspelled key (say `project_idd`) loads, but the engine warns — and when there is no project to mount, the Desktop App's status names the unrecognized key(s), because a typo like that is otherwise invisible: the setting simply never applies. The key list is derived from the engine's own parser, so the table above stays current with it.
- **A UTF-8 byte-order mark (BOM) is tolerated.** Windows Notepad and PowerShell 5.1's `-Encoding UTF8` prepend one; the engine strips it. Still, prefer a plain UTF-8 editor.
- **Environment variables override the file.** If a launcher (systemd unit, Docker `-e` , a scheduled task) sets `SWARMFILE_HUB_URL` , editing `hub_url` in the file will appear to do nothing. Check what actually took effect with `--print-env-vars` (below).
- **Don't fight the Desktop App.** For keys the Desktop App or `swarmfile` CLI manage (mount point, seed mode, hub-only, cache size, read-ahead, throttle, LAN placement, branch), prefer their controls — the Desktop App may rewrite the file and overwrite a manual edit. Hand-editing is for machines with no Desktop App.

Connecting to the hosted service

A released build ships pointing at Swarmfile's production service, so you do **not** need to configure the hub or identity provider by hand. With nothing set, the engine defaults to:

Setting	Default
<code>hub_url</code>	<code>https://hub.swarmfile.com</code>
<code>oidc_issuer</code>	<code>https://id.swarmfile.com</code>
<code>oidc_client_id</code>	<code>swarmfile-desktop</code> on a packaged install; <code>swarmfile-engine</code> is the bare binary's fallback when no bundled config supplies one

A headless machine on the hosted service therefore only needs to say *which* org and project it serves and supply a credential — `org_id` , `project_id` , and one of `SWARMFILE_API_KEY` or `SWARMFILE_OIDC_REFRESH_TOKEN` (see below). Set `hub_url` , `oidc_issuer` , or `oidc_client_id` only when you run a **self-hosted or Enterprise hub** on your own domain.

Keys you can set in the file

Every key is optional. Booleans are JSON booleans (`true` / `false`) in the file — the `true` / `1` string form is only for the environment-variable equivalents.

Key	Type	Default	Env override	What it does
hub_url	string	https://hub.swarmfile.com	SWARMFILE_HUB_URL	Hub API base URL
org_id	string	(none)	SWARMFILE_ORG_ID	Multi-tenant organization ID
project_id	string	(none)	SWARMFILE_PROJECT_ID	Scope every headless/service
branch	string	the project's default branch (each project names it at creation — commonly main)	SWARMFILE_BRANCH	Branch this project uses
office_id	string	default	SWARMFILE_OFFICE_ID	Peer-discovery set <office_id> for discovery in requires a remote
oidc_issuer	string	https://id.swarmfile.com	SWARMFILE_OIDC_ISSUER	OIDC issuer URL
oidc_client_id	string	swarmfile-desktop on a packaged install	SWARMFILE_OIDC_CLIENT_ID	OIDC client ID to swarmfile appear to discover
site_url	string	(install config)	SWARMFILE_SITE_URL	Public site to share URLs (the installer note below)
updater_endpoint	string	(install config)	SWARMFILE_UPDATER_ENDPOINT	Update-manager installer's remote endpoint (note below)
encryption_project_id	string	falls back to project_id when encryption is enabled and no SWARMFILE_ENCRYPTION_KEY is supplied	SWARMFILE_ENCRYPTION_PROJECT_ID	Project who's key is used for encryption
read_only	bool	false	SWARMFILE_READ_ONLY	Force a read-only mode
session_kind	owner member guest	owner	SWARMFILE_SESSION_KIND	Active session kind (badge).
hub_only	bool	false	SWARMFILE_HUB_ONLY	Disable all F from the hub
seed_mode	bool	false	SWARMFILE_SEED_MODE	Run as a NA to cloud storage
ec_lan_placement	bool	false	SWARMFILE_EC_LAN_PLACEMENT	Deliberately place on LAN
lan_from_office	bool	false	SWARMFILE_LAN_FROM_OFFICE	Treat same-unavailable. block mDNS nothing, so to spread across members via the same above.
node_auth	bool	false	SWARMFILE_NODE_AUTH	For orgs that the hub can authoritative to use until enforces no through this
sync_ignore_enabled	bool	true	SWARMFILE_SYNC_IGNORE_ENABLED	Whether .git sync (the git rules" escape listings, excludes every
iroh_bind_addr	string	(OS-assigned ephemeral)	SWARMFILE_IROH_BIND_ADDR	Fixed UDP bind address Defaults to :iroh firewall req. settable, inc
throttle_enabled	bool	false	SWARMFILE_THROTTLE_ENABLED	Enable bandwidth explicit site so set site_down to the real limit
site_download_bandwidth_mbps	number	(unset)	SWARMFILE_SITE_DOWNLOAD_BANDWIDTH_MBPS	Download bandwidth limit
site_upload_bandwidth_mbps	number	(unset)	SWARMFILE_SITE_UPLOAD_BANDWIDTH_MBPS	Upload bandwidth limit
office_hours_start	string	08:00	SWARMFILE_OFFICE_HOURS_START	Office-hour start time to office_hours_end

Key	Type	Default	Env override	What it does
				settable per hours-sta
office_hours_end	string	18:00	SWARMFILE_OFFICE_HOURS_END	Office-hour office_ho hours windc
office_hours_wan_pct	number	20	SWARMFILE_OFFICE_HOURS_WAN_PCT	WAN bandw
off_hours_wan_pct	number	80	SWARMFILE_OFF_HOURS_WAN_PCT	WAN bandw
cache_max_bytes	number	10737418240 (10 GiB)	SWARMFILE_CACHE_MAX_BYTES	Local block-
readahead_buffer_secs	number null	(unset — each file type's own default; 15 s for video containers)	SWARMFILE_READAHEAD_BUFFER_SECS	Read-ahead playback. A file type's d containers, Settings →
changeset_idle_secs	number	300	SWARMFILE_CHANGESET_IDLE_SECS	Idle second: disables grc
attr_cache_secs	number	2	SWARMFILE_ATTR_CACHE_SECS	macOS and Windows/W 5–60s defai
mergetool_cmd	string	(unset)	SWARMFILE_MERGETOOL_CMD	Command to tool runs, the same \$ convention swarmfile u: swarmfile fa tool — any nothing here
mergetool_name	string	(unset)	SWARMFILE_MERGETOOL_NAME	A cosmetic in messages; also set; it c git config.
data_namespace	string	(unset)	SWARMFILE_DATA_NAMESPACE	Per-install d production) something y selection ke
auto_open_changelist_for_preference	bool	false	SWARMFILE_AUTO_OPEN_CHANGELIST_FOR_PREFERENCE	Commit-mo changelist," have to rem overrides ar switch. Off l open chang
deferred_namespace	bool	true	SWARMFILE_DEFERRED_RENAME (alias SWARMFILE_DEFERRED_NAMESPACE)	Apply non-c deletes of c confirmation false to c env value is env spelling
mount_point	string	~/Swarmfile (Unix), S: (Windows)	SWARMFILE_MOUNT_POINT	Where to m on Windows configure th opening adc
mount_point_auto	(managed)	—	—	Do not set. selected for a letter the i

Additional mounts aren't configured here. One engine process can hold several full mount points open at once against the same project/org — see [swarmfile mounts](#) — but that's an entirely runtime, CLI/Desktop-App-driven operation with no config-file or env-var equivalent: there's no `mount_point_2` key or similar, and no way to declare a second mount up front for the engine to open at boot. `mount_point` / `mount_point_auto` above describe the boot mount specifically, whether or not anything else is opened later in the session. You don't need to edit this file to get a free drive letter for an additional mount: `swarmfile mounts open auto` (or the Desktop App's Open mount form, which defaults to it on Windows) picks one.

A note on environment-selection keys

`hub_url`, `oidc_issuer`, `oidc_client_id`, `site_url`, `updater_endpoint`, and `data_namespace` are handled differently from every other key in the table above. A packaged install (the macOS `.app`, the Windows installer, the Linux `.deb`) ships its own read-only install config that sets these — that's how an install stays on the environment it was packaged for, and why one signed binary set can serve different environments without drifting. When the install config sets one of these six keys, **it wins over whatever you put in your per-user config.json** — a hand-edit of `hub_url` or `oidc_issuer` at the per-user path documented above will silently appear to do nothing on a normal packaged install. Your per-user value is only used as a fallback for a key the install config left unset — which is the normal case for a from-source/dev build, or an older single-file-layout install. Every other key in this file follows the ordinary env-var-then-file precedence at the top of this page; only these six are install-config-first.

Settings that are environment-variable-only

Many settings can be tuned via `SWARMFILE_*` env vars but are **not** readable from `config.json` — the file supports only the keys in the table above. [Environment Variable Index](#) is the full list; the notable ones include:

- **Credentials:** `SWARMFILE_OIDC_REFRESH_TOKEN` (non-interactive OIDC), `SWARMFILE_API_KEY` (project-scoped headless key — takes precedence over a refresh token), `SWARMFILE_ENCRYPTION_KEY`, `SWARMFILE_SWARM_KEY`. Credentials are deliberately kept out of the file.
- **Paths & identity:** `SWARMFILE_CACHE_DIR`, `SWARMFILE_IROH_DIR`, `SWARMFILE_USER_ID`, `SWARMFILE_MACHINE_ID`.
- **Encryption toggle/mode:** `SWARMFILE_ENCRYPTION_ENABLED`, `SWARMFILE_EC_ENABLED` (on/off/auto).
- **Networking:** `SWARMFILE_IROH_PUBLIC_ADDR`, `SWARMFILE_HEALTH_PORT`. `SWARMFILE_MDNS_INTERFACES` (comma-separated interface names) restricts LAN discovery to exactly those interfaces and skips the built-in VPN/virtual-interface filter. `SWARMFILE_MDNS_EXCLUDE_INTERFACES` removes more interfaces from it. See [Network Requirements](#). (`SWARMFILE_IROH_BIND_ADDR` / `iroh_bind_addr` is file-settable — see the table above.)
- **QoS:** `SWARMFILE_DSCP_INTERACTIVE`, `SWARMFILE_DSCP_BACKGROUND`, `SWARMFILE_SITE_DOWNLOAD_BANDWIDTH_MBPS` / `SWARMFILE_SITE_UPLOAD_BANDWIDTH_MBPS`. The office-hours schedule keys (`SWARMFILE_OFFICE_HOURS_START` / `_END`, `SWARMFILE_OFFICE_HOURS_WAN_PCT`, `SWARMFILE_OFF_HOURS_WAN_PCT`) are file-settable — see `office_hours_*` / `off_hours_wan_pct` in the table above.
- **LAN shard placement tuning:** `SWARMFILE_EC_LAN_PLACEMENT_REPLICATION` (default 3) is the replication factor K — each erasure shard is deliberately assigned to this many LAN peers. `SWARMFILE_EC_LAN_PLACEMENT_POLL_INTERVAL` (default 300) is the reconciler's safety-net pass interval in seconds, run alongside the reactive pass that fires on a peer-roster change. Both only matter when `ec_lan_placement` is on.
- **Seed tuning:** `SWARMFILE_SEED_POLL_INTERVAL` (default 60, seconds), `SWARMFILE_SEED_UPLOAD_CONCURRENCY` (default 8, concurrent cloud uploads).
- **Runner (headless CI):** `SWARMFILE_RUNNER_MODE` — watches `SWARMFILE_BRANCH` for commits/tags matching a rule in `.swarmfile/runner.yml` and runs it. Env-only by design, no `config.json` key or Desktop App toggle — see [Runner \(Headless CI\)](#). Mutually exclusive with `seed_mode`.
- **Streaming an upload in progress:** `SWARMFILE_STREAM_BUFFER_SECS`, `SWARMFILE_STREAM_READ_WAIT_SECS`, `SWARMFILE_STREAM_WAIT_CONCURRENCY` — env-only, no `config.json` key. `STREAM_BUFFER_SECS` (default 30) is how many **seconds** of playback a `fetch --follow` job keeps buffered ahead of the reader; it is seconds rather than bytes because 30 s of a proxy and 30 s of a full-resolution master are wildly different byte counts and one figure would be wrong for both. `STREAM_READ_WAIT_SECS` (default 15) is how long a read of a not-yet-uploaded range waits before giving up — see below. `STREAM_WAIT_CONCURRENCY` (default 4, clamped 1–64) caps how many reads may be *waiting* at once, so a burst returns the retryable "try again" promptly instead of parking every operation thread; it is sized when the mount starts, so a change needs a remount.
- **Peer-to-peer budgets:** `SWARMFILE_IROH_PEER_BUDGET_SECS` (default 30) and `SWARMFILE_IROH_PEER_BUDGET_BACKGROUND_SECS` (default 10) cap how long a peer connection attempt may take for interactive and background work respectively.
- **Durable create queue tuning:** `SWARMFILE_CREATE_DISPATCH_CONCURRENCY` (default 8, 1–64 — creates sent to the hub at once), `SWARMFILE_PENDING_CREATE_WAIT_SECS` (default 8, 0–120 — how long a lock/ACL/share/comment on a not-yet-confirmed file waits before answering `409 pending_create`), and the coalescing/batching knobs (`SWARMFILE_CREATE_COALESCE`, `_HOLD_SECS`, `_HOLD_UPLOADING_SECS`, `_BATCH`). See [Coding Agents](#).
- **Worksharing refusal mode:** `SWARMFILE_WORKSHARING_FAIL_CLOSED` (Windows only) makes a native worksharing open refuse when the hub can't be reached.
- **mdNS service type:** `SWARMFILE_MDNS_SERVICE_TYPE` overrides the `_swarmfile._udp.local.` service type — for isolated test networks, not for normal use.
- **Upload diagnostics:** `SWARMFILE_UPLOAD_TRACE=1` promotes per-attempt upload logging to `warn`, including the hub's status and response body when a block upload or commit is refused — the first thing to set when a save isn't landing and you need to know why. Off by default; the lines are noisy for ordinary work.
- **Misc:** `SWARMFILE_XREF_PREFETCH`, `SWARMFILE_DOCTOR_PERIOD_SECS`, `SWARMFILE_OTEL_ENDPOINT`, `SWARMFILE_OTEL_SERVICE_NAME`.

How streaming an in-progress upload behaves

Always on: a file that has never finished uploading can be opened through the mount — it appears at its full eventual size and reads work, waiting briefly where the bytes have not landed. Worth knowing how it behaves, because a read into a range that has not been uploaded yet has to wait, and this runs on a filesystem operation thread. An operation that waits too long does not stall one read — it takes the whole drive with it. The wait is therefore strictly bounded (`SWARMFILE_STREAM_READ_WAIT_SECS`, default 15 s, and `0` does not mean "forever" — it falls back to the default), and a read that outruns the upload returns a retryable "not yet" rather than an I/O error, because an application that sees an I/O error part way through a file will usually conclude the file is damaged and discard your document.

Different NLEs and DCC tools react differently to a delayed read; the bound is tuned to release the drive before an editing application gives up on it. Two consequences worth knowing: only so many reads may be *waiting* at once (a burst beyond that gets the "try again" immediately rather than parking every operation thread), and only a file with **no** committed version streams — a re-save of an existing file is readable once it commits, never mid-save.

Scope worth knowing:

- It applies **only to files with no committed version yet** — a new file landing for the first time. A file you are already reading never changes size or content because somebody started a new save.
- A file that *is* streaming reports the size it is going to be, not `0`. Without that, nothing could read it: no application asks for bytes past the end of an empty file.
- If the uploading machine is on the same **LAN**, it serves those blocks directly from its own disk, before they reach the cloud. Over a WAN it does not — that would carry every block twice on the same contended uplink.

And the limit that no setting changes: this decides **which** bytes arrive first, not how fast the link is. A 4 TB file on a 100 Mbit/s uplink delivers about 12 MB/s however well it is ordered.

Project-local files (`.swarmfile/`)

Some configuration lives **inside a project** — committed like any other file, so it versions with the branch and every contributor gets the same behavior — instead of in the engine's `config.json`:

File	Keys	What it does
<code>.swarmfile/runner.yml</code>	rule list	Headless CI rules — see Runner (Headless CI) .
<code>.swarmfile/merge.yml</code>	<code>diff3: true false</code>	Project default for <code>swarmfile merge --diff3</code> and the tray's "Auto-resolve conflicts": on a merge conflict, try a three-way text merge of each conflicting file before giving up. <code>true</code> turns it on by default; <code>false</code> turns it off and the CLI/tray name this file as the reason. Absent means off (auto-resolve stays opt-in). An explicit <code>--diff3 / --no-diff3</code> always wins over the file.
<code>.swarmfile/cache.yml</code>	<code>caches: list of name, env, subdir</code>	Tool-cache redirects for <code>swarmfile run</code> : each entry sets the <code>env</code> variable a tool reads for its cache location (e.g. <code>PNPM_HOME</code> , <code>CARGO_HOME</code>) to this project's shared machine-local scratch directory plus <code>subdir</code> . Every mount of the project on this machine shares those directories, so a second mount does not re-download what the first already fetched. A bad entry is dropped with a warning, never the whole file. See scratch .
<code>.swarmfile/publicignore.yml</code>	<code>version, paths, derived.ci_logs</code>	Path globs (gitignore-style, with <code>!</code> re-includes) kept out of a release's public view; <code>derived.ci_logs</code> sets CI-log exposure (<code>member-only</code> by default). Read from the tree of the tag being published and enforced as the release is built. See Publishing Releases .

The runner, merge and cache files are read from the branch the operation runs against, so they can differ per branch; `.swarmfile/publicignore.yml` is read from the tree of the tag being published. A `.swarmfile/merge.yml` that is missing or unparseable is treated exactly like `diff3: false` — never as "turn extra automatic behavior on".

For example, to make text-conflict auto-resolve the default for a project:

```
# .swarmfile/merge.yml
diff3: true
```

Engine command-line flags

Beyond reading `config.json`, the `swarmfile-engine` binary takes a handful of one-shot flags — useful on a headless node where there's no Desktop App to click.

Answered immediately — about this binary, or a running engine:

Flag	What it does
<code>--status</code>	Engine, hub, peers, and pending work at a glance. Queries a running engine; falls back to a one-shot read of local state if none is running.
<code>--print-env-vars</code>	The resolved <code>SWARMFILE_*</code> values this build consumed, secrets redacted, with a footer naming the install-config and per-user config files that fed the merge (or the <code>SWARMFILE_CONFIG</code> path).
<code>--capabilities</code>	The compiled feature set of <i>this</i> binary — which optional features it was built with.
<code>--build-info</code>	The same compiled features plus the target architecture, in a one-line machine-readable form, then exit.
<code>--version</code> , <code>-V</code>	Print the version and exit.
<code>--help</code> , <code>-h</code>	Print the flag list and exit.

Lifecycle — against a running engine:

Flag	What it does
<code>--pause</code>	Pause sync without unmounting the drive.
<code>--resume</code>	Resume after <code>--pause</code> .
<code>--quit</code>	Ask the running engine to drain in-flight work and stop.

`--help` and `--version` answer *before* the engine reads any config, so they work on a broken or unconfigured install. For day-to-day version-control and mount operations, use the [swarmfile](#) CLI — it drives the same engine over its control socket — rather than these flags.

The engine's own `--help` also lists offline/pinning flags (`--pin`, `--prepare-offline`, and related). These currently run only while the engine is **stopped**, so they can't be driven against a live mount and aren't part of the normal headless workflow. To force a scope resident for offline or render-farm use against a *running* engine, use [swarmfile hydrate](#) instead.

Verifying what actually took effect

Because several sources feed the final config, print the **resolved** values the engine will use — with secrets redacted, and a footer naming the install-config and per-user files that fed the merge (or the `SWARMFILE_CONFIG` path):

```
swarmfile-engine --print-env-vars
```

`swarmfile doctor` (or the standalone `swarmfile-doctor`) then confirms the resolved hub URL, org, and credentials actually reach the hub. See [CLI: swarmfile-doctor](#).

Example: a headless seed node

On the hosted service the hub and IdP are already the defaults, so a seed node only names its org and project:

```
{
  "org_id": "org_abc123",
  "project_id": "proj_xyz789",
  "seed_mode": true,
  "cache_max_bytes": 107374182400
}
```

(On a **self-hosted or Enterprise** hub, add `hub_url`, `oidc_issuer`, and `oidc_client_id` for your own domain.)

Pair this with a credential in the environment — `SWARMFILE_API_KEY=sf_key_...` for a project-scoped key, or `SWARMFILE_OIDC_REFRESH_TOKEN=...` — since credentials aren't stored in the file. See [Deployment Topologies](#) and [Self-Hosted Seed Nodes](#) for the full headless-node walkthrough, [Network Requirements](#) for exactly what a machine using this file needs to reach, and [Deploying to Your Team](#) for pre-staging this file across a fleet.

Network Requirements

What Swarmfile actually talks to on the network, so you can configure a firewall allowlist before rollout instead of discovering it by trial and error. Everything below is what the desktop client (engine + Desktop App) initiates outbound — nothing here requires an inbound port opened on your network firewall for ordinary use. Host firewalls on each machine are a separate matter: LAN peer discovery needs inbound UDP from the local subnet, which the Windows installer allows for you (see [LAN peer discovery](#)).

Summary

Purpose	Protocol	Port	Host(s)	Required?
Control plane — metadata, auth, commits, presigned URLs, live sync	HTTPS	443	<code>hub.swarmfile.com</code> , or your own hub if self-hosted	Always
Identity — sign-in, token refresh	HTTPS	443	<code>id.swarmfile.com</code> , or your own IdP if you bring your own SSO	Always
File content — block upload/download	HTTPS	443	the org's storage origin: <code>*.r2.cloudflarestorage.com</code> by default, or the org's dedicated/bucket endpoint (a customer S3-compatible endpoint under BYOS)	Always — this is the actual file data; the hub only mints short-lived presigned URLs
Installer downloads	HTTPS	443	<code>releases.swarmfile.com</code>	Install and repair only
git-LFS agent installer	HTTPS	443	<code>get.swarmfile.com</code>	Only while installing/updating the git-LFS agent
Auto-update check	HTTPS	443	<code>updates.swarmfile.com</code>	Whenever the Desktop App is running
Peer-to-peer file transfer	QUIC (UDP)	OS-assigned by default	direct to other peers' IPs	Only when P2P is enabled (the default; off under cloud-only mode)
P2P relay and peer discovery (NAT traversal)	HTTPS/443 + UDP/7842	third-party hosts — see below	Only when P2P is enabled	
LAN peer discovery	mDNS, UDP 5353 multicast	local subnet	Only when P2P is enabled	
Dashboard / web app	HTTPS	443	<code>swarmfile.com</code>	User-initiated browsing only, opened in your default browser — not the engine or Desktop App process

Which exact hub/identity/releases/updates hostnames a given install uses is a per-install `config.json` setting, not something baked into the binary at compile time — a fresh install on the standard hosted plan defaults to the prod hosts above; check that file on a specific machine if you need to confirm. See [Engine Config File](#).

Control plane and identity

The engine talks to one hub host for control-plane work: browsing and editing metadata, commits, minting the short-lived presigned URLs for block transfer, and a live WebSocket (`wss://`, same host) that pushes changes from other users in real time. **File content does not flow through the hub.** For a read the hub answers with a redirect to a presigned URL and the client fetches the block bytes directly from the org's storage origin — `*.r2.cloudflarestorage.com` on the hosted default, a jurisdiction-pinned endpoint, or a customer-supplied S3-compatible endpoint under [BYOS](#); uploads are presigned the same way. Sign-in and token refresh go to a separate identity host — the built-in IdP by default, or your own OIDC provider if you've configured bring-your-own SSO (see [Identity](#)).

Both are plain HTTPS on port 443. No other port is used for control-plane or identity traffic.

git-LFS

On a project used as a git-LFS server, `git lfs` talks to the hub host for the Batch API, locks, and the transfer endpoint — and, for the stock `basic` transfer (and any push ≥ 64 MiB through the desktop agent), to the storage origin directly via a presigned URL, same as block traffic. Downloads stream through the hub. The git-LFS agent installer is fetched once from `get.swarmfile.com` (see the summary table). Everything stays on 443.

Peer-to-peer file transfer

By default, machines in the same office exchange file data directly over QUIC (UDP) rather than always routing through the cloud — this is what makes LAN-first fast. The engine binds to an OS-assigned ephemeral UDP port for this; it does not listen on a fixed, predictable port unless you explicitly set one (`"iroh_bind_addr"` in `config.json`, or `SWARMFILE_IROH_BIND_ADDR` in the environment — an operator choice for a specific deployment, not a default).

Peer-to-peer relay and discovery

Two pieces of this are third-party infrastructure Swarmfile doesn't operate itself, run by Number Zero (the maintainers of the P2P transport Swarmfile builds on):

- **Relay fallback**, used when two peers can't establish a direct connection (e.g. both behind restrictive NATs): a handful of regional relay hosts (`*.relay.n0.iroh.link`), reached over HTTPS/443 with QUIC address-discovery on UDP/7842.
- **Peer discovery**, used to look up how to reach another peer: `dns.iroh.link`, over HTTPS and plain DNS.

If your firewall policy needs to name every host P2P might touch, these are worth listing alongside your own hub domain — though in practice, if UDP/QUIC is blocked entirely (common on locked-down corporate networks), the simpler answer is usually [cloud-only mode](#) rather than trying to allowlist P2P traffic through it.

LAN peer discovery

Finding peers on the same local network uses standard mDNS (multicast DNS, the same mechanism as Bonjour/AirPlay discovery) — no custom protocol, no non-standard port. If your network already blocks mDNS multicast (some corporate/VLAN setups, and Docker bridges, do this by default), Swarmfile still works — it just can't discover LAN peers automatically, and every read falls back to the peer-relay path or the cloud. There's a separate `lan_from_office` setting for networks that block multicast but where you still want deliberate LAN peer placement; see [Deployment Topologies](#).

mDNS runs only on physical LAN interfaces (Ethernet, Wi-Fi). The engine skips VPN and tunnel adapters (macOS `utun*`, WireGuard, OpenVPN, Tailscale, ZeroTier), Hyper-V/WSL, VMware, VirtualBox, Parallels and Docker virtual networks, and Apple's AirDrop links (`awdl0` , `llw0`), because none of them reach a colleague's machine on the office LAN. It advertises only its LAN addresses, never a VPN address. It checks interfaces again every 30 seconds, so connecting a VPN later doesn't move discovery onto it. To override the choice, set `SWARMFILE_MDNS_INTERFACES` (use only these interfaces, comma-separated) or `SWARMFILE_MDNS_EXCLUDE_INTERFACES` (never use these). `swarmfile doctor` 's `mdns listen` check lists the interfaces in use and those it skipped. It warns when no LAN interface is available.

Windows Firewall. Windows filters *inbound* traffic per program. If `swarmfile-engine.exe` has no inbound rule, other machines' mDNS queries and announcements never reach it, but its own outgoing packets still leave. The result is one-sided: other machines may briefly see this one, and this one never finds anyone. Windows only offers to create a rule when someone runs the engine interactively, and then only for the Public profile, so silent/GPO installs and standard users would otherwise get nothing.

The installer (both the bundled `Setup.exe` and the bare `.msi`) therefore adds one inbound rule, and removes it on uninstall:

Setting	Value
Display name	Swarmfile – local network file sharing
Group	Swarmfile
Program	C:\Program Files\Swarmfile\swarmfile-engine.exe
Direction / action	Inbound, Allow
Protocol / ports	UDP, any local port (mDNS on 5353, plus the QUIC port, which is ephemeral unless you pin it with <code>iroh_bind_addr</code>)
Remote addresses	Local subnet only
Profiles	Domain, Private (not Public)

A second rule in the same group, `Swarmfile – local network check (doctor)` , does the same for `swarmfile-doctor.exe` , so its one-second `mdns listen` check can hear replies when no engine is running. To include the Public profile as well, install with `msiexec /i Swarmfile-x64.msi SWARMFILE_FW_PROFILES=7 /qn` (the value is a bitmask: 1 = Domain, 2 = Private, 4 = Public). Check the result with:

```
Get-NetFirewallRule -Group Swarmfile | Format-Table DisplayName, Enabled, Profile, Direction, Action
Get-NetFirewallRule -Group Swarmfile | Get-NetFirewallAddressFilter # RemoteAddress: LocalSubnet
```

If you deploy from the `.zip` , run the engine from a different path, or your GPO sets **Apply local firewall rules: No** (which ignores rules an installer adds), create the equivalent rule yourself, through GPO/Intune or with:

```
New-NetFirewallRule -DisplayName "Swarmfile – local network file sharing" -Group "Swarmfile" `
-Direction Inbound -Action Allow -Protocol UDP `
-Program "C:\Program Files\Swarmfile\swarmfile-engine.exe" `
-RemoteAddress LocalSubnet -Profile Domain,Private
```

To avoid a program-scoped rule, pin the QUIC port (`"iroh_bind_addr": "0.0.0.0:4433"` in [config.json](#)) and open UDP 5353 and 4433 from `LocalSubnet` with `-LocalPort 5353,4433` instead of `-Program` . A **Block** rule for `swarmfile-engine.exe` (created when someone clicked Cancel on a firewall prompt) overrides any allow rule, so delete it if one exists. `swarmfile doctor` 's `mdns listen` check warns when the hub lists peers in your office or on your subnet but mDNS has found none of them. That usually means a firewall or a network that filters multicast.

Linux firewalls

The `.deb` doesn't change your firewall. On a host running `ufw` or `firewalld` with a default-deny inbound policy, allow UDP 5353 (mDNS) and the engine's QUIC traffic from the local subnet. QUIC uses a random UDP port unless you pin it with `iroh_bind_addr` , so you can either allow all UDP from the subnet or pin the port and open only that one. The examples use `192.168.1.0/24` . Use your own subnet, which `ip -4 addr` shows (an address of `192.168.1.23/24` is on `192.168.1.0/24`).

The tidy option is to pin the port. Add `"iroh_bind_addr": "0.0.0.0:4433"` to [config.json](#) , restart the engine (`systemctl --user restart swarmfile-engine`), and then:

```
# ufw
sudo ufw allow proto udp from 192.168.1.0/24 to any port 5353 comment 'Swarmfile mDNS'
sudo ufw allow proto udp from 192.168.1.0/24 to any port 4433 comment 'Swarmfile QUIC'

# firewalld
sudo firewall-cmd --permanent --add-rich-rule='rule family="ipv4" source address="192.168.1.0/24" service name="mdns" accept'
sudo firewall-cmd --permanent --add-rich-rule='rule family="ipv4" source address="192.168.1.0/24" port port="4433" protocol="udp" accept'
sudo firewall-cmd --reload
```

If you leave the port unpinned, allow all UDP from the subnet instead:

```
# ufw
sudo ufw allow proto udp from 192.168.1.0/24 to any comment 'Swarmfile LAN'

# firewallld
sudo firewall-cmd --permanent --add-rich-rule='rule family="ipv4" source address="192.168.1.0/24" protocol value="udp" accept'
sudo firewall-cmd --reload
```

Some stock rule sets already allow part of this. Ubuntu's default ufw rules accept multicast mDNS, and Fedora Workstation's default firewall zone allows mDNS and UDP ports above 1024. The explicit rules do no harm there. They also cover the unicast replies that some mDNS responders send. If your LAN uses IPv6 only, add the same rules for your IPv6 prefix.

`swarmfile doctor`'s `host firewall` check reads these firewalls without root and warns when one is active with no rule that allows UDP 5353 and the QUIC port. `ufw` only shows its rules to root, so as a normal user the check can't confirm them. In that case it warns only when `mdns listen` also reports missing LAN peers.

macOS

The macOS Application Firewall allows the signed engine by default. Two settings stop LAN peers from reaching it: **Block all incoming connections** (System Settings → Network → Firewall → Options...), and a *Block incoming connections* entry for `swarmfile-engine`, which appears when someone clicks Deny on the firewall prompt. The `host firewall` check reports both and gives the command that reverses each one.

On macOS 15 and later, **Local Network privacy** also gates LAN access. The Desktop App asks for it, and the engine installed inside the app is covered by that permission (System Settings → Privacy & Security → Local Network). An engine binary running outside the app bundle, such as a render-farm or CI node, has no such permission. Unless it runs as root (a LaunchDaemon) or in the foreground of a Terminal or SSH session that stays open, macOS can deny it LAN access without telling anyone. Connections to LAN peers or a LAN hub then fail with "No route to host", and mDNS finds nobody, while internet traffic keeps working. The `macos local network` check warns when the engine runs this way. Apple's [TN3179: Understanding local network privacy](#) lists these rules.

Cloud-only mode (hub-only)

For a network that blocks UDP/QUIC outright, or a security policy that wants zero peer-to-peer connections and zero peer IP exposure, **cloud-only mode** turns P2P off entirely — not just "prefers not to use it." (The engine and config file call it `hub_only`; the Desktop App and CLI label the same switch *Cloud-only mode*.) With it enabled, the engine never opens a QUIC endpoint, never starts mDNS, and never dials another peer. The entire network surface collapses to the HTTPS/443 rows in the summary table above — metadata to the hub, file bytes to the org's storage origin (control plane, identity, storage, and — while the Desktop App is active — installer/update checks), so a firewall must allow both, not just the hub. Nothing UDP is ever touched.

This is set per-machine (`"hub_only": true` in `config.json`, or `SWARMFILE_HUB_ONLY=1` in the environment), or centrally for the whole org by an owner or admin (Settings → Network). The two combine with OR, not override: an org-enforced policy can only ever turn cloud-only **on** for a machine, never force it back **off** — a machine an operator has locally pinned to cloud-only stays cloud-only regardless of what the org policy says. See [Security](#) for how this fits the broader threat model, and [Operations](#) for where the setting shows up in your audit log.

Verifying what a specific machine actually needs

`swarmfile doctor` runs a battery of connectivity checks — DNS resolution, HTTPS reachability to the hub, an authenticated round-trip, a QUIC dial test (skipped cleanly when cloud-only mode is on), and an mDNS listen test (same) — and reports pass/warn/fail for each, with a `--json` mode for pasting into a firewall-exception ticket. Run it from a machine on the network you're configuring rather than guessing; see [swarmfile-doctor](#).

What's not covered here

Self-hosted seed and NAS nodes participate in the same P2P swarm as any other machine — no additional inbound ports are required for them either. A seed node run as a headless cloud service (rather than on-prem hardware) can optionally expose a health-check HTTP endpoint for your own orchestration (Kubernetes, Fly.io, etc.) — that's an explicit opt-in for that specific deployment shape, not something a normal desktop or on-prem seed install does, and it's the one case where the engine listens on all interfaces rather than loopback-only. Even then, only `/livez` (a bare `{"status":"ok"}`) answers a non-loopback caller — `/healthz`'s detailed body, which exposes sensitive operational detail, 403s for anything but loopback regardless of which interface the port is bound to, so a container orchestrator's liveness probe works but a remote caller can't use the same port to enumerate your fleet. Talk to us if you're deploying that way and need the specifics.

Storage Format


This page specifies how Swarmfile stores a file's bytes in object storage: how a file is split into blocks, how blocks are named, how the list of blocks (the *manifest*) is encoded, how blocks are encrypted, and where they land in a bucket. It is written so that someone holding the stored objects (and, for an encrypted project, the project key) can reconstruct a file with no Swarmfile software involved. It describes the format the current code writes and reads. Where something isn't pinned down, this page says so rather than guessing.

It covers **file content only**. File names, folder structure, versions, branches, commits, ACLs, and comments live in the project's metadata on the hub, not inside blocks, and their formats aren't specified here. In practice you need that metadata to know which manifest belongs to which path: each file entry records its manifest's identifier as `rootCid`.

If you only want your files back, you don't need this page: copy them off the mount, or use [Branch Mirror](#) to export real files to your own bucket. See [Data Portability & Offboarding](#).

Overview

- A file is split into **chunks**, usually with content-defined chunking (FastCDC) averaging 1 MiB.
- Each chunk is stored as one **block**, named by its **CID**: the BLAKE3 hash of the chunk's *plaintext* bytes.
- A **manifest**, a small JSON document listing the chunks in file order, is itself stored as a block with its own CID. That CID is the file entry's `rootCid`.
- A very large file's manifest is split into **segment** blocks under a small **segmented root** (manifest format v2).
- On encrypted projects, every block, manifests included, is stored as AES-256-GCM ciphertext under a key derived from the project key and the block's CID. The CID is still the plaintext hash.
- Some chunks also get **Reed-Solomon 10+4 erasure-coded shards**, stored as extra blocks *in addition to* the chunk itself.



Chunking

Default parameters

Constant	Value
FastCDC minimum chunk	262,144 bytes (256 KiB)
FastCDC average (target) chunk	1,048,576 bytes (1 MiB)
FastCDC maximum chunk	4,194,304 bytes (4 MiB)
Fixed-size chunk (when fixed chunking is used)	1,048,576 bytes (1 MiB)

The CDC implementation is FastCDC "v2020" from the Rust `fastcdc` crate (3.2.x), with its default normalization (level 1) and gear table. A compatible implementation in the hub uses the same parameters, so CIDs dedupe across both.

You never need to re-run the chunker to read a file. The manifest records every chunk's offset and size, so a reader just follows the manifest. The chunking parameters only matter if you want to reproduce Swarmfile's CIDs for new data.

File-type profiles

On the mount's write path, the chunking strategy depends on the file extension (case-insensitive):

Extensions	Strategy
everything not listed below (incl. <code>.rvt</code> , <code>.dwg</code> , <code>.max</code> , <code>.ifc</code> , <code>.drp</code> , <code>.settings</code> , <code>.prproj</code>)	CDC 256 KiB / 1 MiB / 4 MiB
<code>.pdf</code> , <code>.dwf</code>	Fixed 1 MiB
<code>.slog</code> , <code>.rws</code> , <code>.dat</code> , <code>.xml</code> , <code>.avb</code> , <code>.avp</code> , <code>.msn</code> , <code>.pmsm</code> , <code>.nksm</code>	Fixed 1 MiB, plus the small-file fast path

Small-file fast path: a file with one of the fast-path extensions that is smaller than 50 KiB (51,200 bytes) is stored as a **single chunk**, with the strategy recorded as `fixed` and `chunk_size` equal to the file's size (minimum 1).

The profiles also set read-ahead behavior. That doesn't affect the stored format.

Other writers can pick other strategies. For example, `swarmfile-seed` uses fixed-size chunks unless you pass `--cdc`, and small browser uploads are written as a single chunk. Whichever strategy was used is recorded in the manifest's `chunking` field. The per-extension table above may change. It is not part of the stable format.

Block identifiers (CIDs)

A CID is the **BLAKE3-256 hash** of the block's **plaintext** bytes, written as **64 lowercase hex characters**. There is no multihash or multibase prefix, no version byte, and no domain separation.

Block kind	CID is BLAKE3 of...
Chunk	the chunk's plaintext bytes
Flat manifest	the manifest's exact serialized JSON bytes
Manifest segment / segmented root	that block's exact serialized JSON bytes
Erasure shard	the shard's bytes as stored (see Erasure coding)

Because a CID hashes plaintext, encrypting a block never changes its CID, and identical content within a project always has one CID. Encryption uses a fresh random nonce every time, so the same plaintext encrypted twice gives different stored bytes. Dedup happens at the CID level, not the ciphertext level.

Verification rule: after you fetch a block (and decrypt it, if it's encrypted), `blake3(plaintext)` must equal the CID you asked for. The only exception is an erasure shard, which is verified against its stored bytes.

Manifests

Manifests are JSON as written by Rust's `serde_json::to_vec`: compact, with no whitespace and fields in the order shown below. Because a manifest's CID is the hash of its exact bytes, **re-serializing a manifest with different formatting or key order gives a different CID**. To read a manifest, any JSON parser works. To reproduce a CID, you have to match the serialization byte for byte.

Flat manifest (format v1)

Used whenever the whole manifest fits in one block, which covers almost every file. It has no version field. Version 1 is implicit.

```
{
  "total_size": 3355443,
  "chunk_size": 1048576,
  "chunks": [
    {"cid": "9f1c...e2a0", "offset": 0, "size": 1210334},
    {"cid": "4b7d...0c19", "offset": 1210334, "size": 862117},
    {"cid": "e03a...77f4", "offset": 2072451, "size": 1282992}
  ],
  "chunking": {"type": "cdc", "min": 262144, "avg": 1048576, "max": 4194304}
}
```

(Formatted and with CIDs shortened for readability. Real CIDs are 64 hex characters, and stored bytes have no whitespace.)

Field	Type	Meaning
<code>total_size</code>	u64	File length in bytes. Equals the sum of every chunk's <code>size</code> .
<code>chunk_size</code>	u32	For <code>fixed</code> : the chunk size. For <code>cdc</code> : the <i>target average</i> , not a real chunk size. Kept for backward compatibility. Use <code>chunking</code> and each chunk's <code>size</code> instead.
<code>chunks</code>	array of ChunkRef	Chunks in file order. An empty file has <code>[]</code> .
<code>chunking</code>	object	How the file was chunked (see below). If absent, readers treat it as <code>{"type":"fixed","chunk_size":1048576}</code> .

`chunking` is an internally tagged object:

- `{"type": "fixed", "chunk_size": N}`
- `{"type": "cdc", "min": N, "avg": N, "max": N}`

ChunkRef:

Field	Type	Meaning
<code>cid</code>	string	CID of the chunk's plaintext (64 lowercase hex).
<code>offset</code>	u64	Byte offset of this chunk in the file.
<code>size</code>	u32	Plaintext length of this chunk.
<code>erasure</code>	object, optional	Erasure-coding group for this chunk (see Erasure coding). Left out when there are no shards.

Chunks are contiguous: each chunk's `offset` equals the previous chunk's `offset` + `size`. The same CID can appear more than once when a file repeats content. Its block is stored once and referenced from each position.

Browser-upload shape (version 1)

Some small files uploaded through the web app may have a manifest in an older envelope, which readers still accept:

```
{"version": 1, "type": "file", "size": 35, "blocks": [{"offset": 0, "size": 35, "cid": "..."}]}
```

`blocks` maps one-to-one onto `chunks`, and `size` onto `total_size` (if `size` is missing, total the block sizes). Treat it as fixed chunking. New writes never produce this shape.

Segmented manifest (format v2)

When a flat manifest would be too big for one block (see [Size thresholds](#)), the chunk list is split into **segment** blocks, and the entry's `rootCid` points at a small **segmented root** instead. A flat manifest's bytes never change when this happens. A file is only segmented if it can't be stored flat.

Segmented root:

```
{
  "manifest_version": 2,
  "total_size": 5497558138880,
  "chunking": {"type": "cdc", "min": 262144, "avg": 1048576, "max": 4194304},
  "segments": [
    {"cid": "a41e...9d03", "first_offset": 0, "chunk_count": 31207},
    {"cid": "07bc...41fe", "first_offset": 32730906112, "chunk_count": 31188}
  ]
}
```

Field	Type	Meaning
<code>manifest_version</code>	u32	Always 2.
<code>total_size</code>	u64	File length in bytes.
<code>chunking</code>	object	Same as in the flat manifest, but always present here.
<code>segments</code>	array of SegmentRef	Segments in file order.

SegmentRef: `cid` (the segment block's CID), `first_offset` (file offset of the segment's first chunk, so a range read can binary-search to the segment it needs), and `chunk_count` (how many ChunkRefs the segment holds).

Segment block:

```
{"chunks": [ {"cid": "...", "offset": 0, "size": 1048576}, ... ]}
```

A segment holds one contiguous slice of the ChunkRef list, in the same shape as the flat manifest's `chunks`, including any `erasure` objects. To rebuild the full chunk list, concatenate the segments' `chunks` in `segments` order. If a segment's length doesn't match its `chunk_count`, treat that as corruption.

There is one level of segmentation. The root is never nested. A file whose segmented root would itself exceed the block limit is refused at write time.

Telling the shapes apart

A root block is parsed like this, in order:

1. Parse as a **flat manifest**, which needs `chunks` and also accepts a `blocks` array. If that works, it's flat.
2. Otherwise parse as a **segmented root**, which needs `manifest_version`, `total_size`, `chunking`, and `segments`. If that works *and* `manifest_version == 2`, it's segmented.
3. Otherwise the block is an error. That includes a segmented root with any `manifest_version` other than 2. Readers fail loudly and never guess.

The two canonical shapes are mutually exclusive: a flat manifest never has `segments` or `manifest_version`, and a segmented root never has `chunks`.

Size thresholds

Limit	Value
Maximum stored size of any manifest, segment, or root block	4,194,368 bytes (4 MiB + 64)
Target <i>plaintext</i> size per segment	¾ of that limit, 3,145,776 bytes

The flat-or-segmented decision is made on the **stored** size: after encryption on an encrypted project, where encryption adds 29 bytes per block. Segments are packed greedily in file order. Each ChunkRef costs its compact-JSON length plus 1 byte, a new segment starts when the next ref would push the open one past the target, and a ref is never split. Readers don't need any of this. It only matters if you want to reproduce the exact segment and root CIDs for a file.

To give a sense of scale: a bare ChunkRef is about 100 bytes, so a flat manifest can hold roughly 40,000 chunks (about 40 GB at the 1 MiB CDC average), and when the list outgrows one block it splits into segments packed to the ¾ target — about 31,000 refs each. A ChunkRef carrying erasure metadata is about 10× larger, so erasure-coded files segment much sooner.

Encryption

Which projects are encrypted

Each project has a fixed **encryption tier**, chosen when it's created:

Tier	Blocks at rest	Who holds the project key
<code>managed</code> (default on paid plans)	AES-256-GCM ciphertext	The hub, which stores the key wrapped by a hub-held key-encryption key (KEK) and hands it over TLS to authenticated clients that are allowed to have it
<code>e2e</code> (opt-in, Pro and above)	AES-256-GCM ciphertext	Only enrolled user devices and, if configured, the org's recovery key. The hub stores only wrapped copies it can't open.
<code>plaintext</code>	Unencrypted	Nobody. There's no key.

The `plaintext` tier is the only tier available on the Free (no-card) plan, and **every public project must use it**. Free-plan and public-project blocks are stored as plain bytes. A paid-plan org that asks for `plaintext` on a private project gets `managed` instead.

What is *not* encrypted, even on the managed tier

- **Directly uploaded git-LFS objects.** An LFS object uploaded by the stock `git-lfs` client through a presigned single PUT, or an object of **64 MiB or more** uploaded by Swarmfile's built-in LFS transfer agent (as a multipart upload), is stored as **one plaintext object** under the `lfs/` keyspace (see [Where blocks are stored](#)). Smaller objects pushed through the agent, and objects that go through the hub's native ingest, become ordinary encrypted blocks with a normal manifest. LFS is refused outright on `e2e` projects. See [Git-LFS](#) and [Security Architecture](#).
- **Some hub-generated derivatives.** The hub renders previews of managed content in memory. Thumbnails, video poster frames and point-cloud preview images are then stored **encrypted** (see [Hub-generated preview images](#)). These stay unencrypted:
 - scrubbable **video proxies** (a 720p H.264 transcode) at `_videoproxy/{rootCid}`, kept up to 30 days;
 - **whole-file copies** the hub assembles to serve browser downloads and previews at `_materialized/{rootCid}`, kept up to 24 hours;
 - short-lived **render inputs and outputs** the preview containers read and write (`_videosrc/`, `_videoposter/`, `_pcpreview/`), deleted when that render finishes;
 - **bare-key preview images:** JPEGs at the bare key `{sha256}`. They aren't rewritten or deleted, and stop being used once the file's content changes.

None of these exist for `e2e` projects, because the hub can't read their content. On the `plaintext` tier, previews are plaintext like everything else.

- **Plaintext managed blocks.** Some managed-tier content was uploaded without envelope encryption; readers detect it by checking whether `blake3(stored bytes) == CID` before trying to decrypt (see below).

Block ciphertext layout

offset	length	field	
0	1	version	0x02 (current), 0x01 (version 1, read-only), or 0x82-0xC0 (a rotated key generation)
1	12	nonce	random, fresh for every encryption
13	N+16	ciphertext	AES-256-GCM(plaintext), with the 16-byte GCM tag appended

Encryption adds exactly **29 bytes** to each block (1 + 12 + 16). There is no associated data (AAD).

The version byte also says **which generation of the project key** sealed the block (see [Key rotation](#)): `0x01` and `0x02` mean generation 1, and `0x80 + g` means generation *g*, for *g* from 2 to 64 (bytes `0x82` to `0xC0`). Every other value is reserved, and readers reject it.

The **per-block key** is derived with HKDF-SHA256 from the 32-byte project key of that generation:

Version byte	HKDF input key	HKDF salt	HKDF info	Output
<code>0x02</code>	generation-1 project key (32 bytes)	the CID as its 64-character lowercase hex ASCII string (64 bytes, not the 32 raw hash bytes)	<code>swarmfile-block-encryption/v1</code> (ASCII)	32 bytes, the AES-256 key
<code>0x82 - 0xC0</code>	the project key of generation <code>version - 0x80</code>	same as <code>0x02</code>	same as <code>0x02</code>	32 bytes
<code>0x01</code> (version 1)	generation-1 project key	none (RFC 5869 default: 32 zero bytes)	the CID's hex ASCII string	32 bytes

The CID used in the derivation is the block's own CID: the chunk's CID for a chunk, and the manifest's, segment's, or root's CID for those blocks. Every block is sealed under a different key, and a block can't be decrypted under the wrong CID.

Compatible implementations exist in the engine, the hub, and the web app. The web app decrypts only E2E share content, with the single project key carried in the share link; an E2E project can be [rotated](#), and a link minted before a rotation carries the pre-rotation key.

How the project key is held (who can decrypt)

This section only explains who can decrypt. It's not needed to parse the format.

- **Managed:** the hub generates a random 32-byte project key and stores it AES-256-GCM-wrapped under its KEK. An authorized, authenticated client gets it over TLS as 64 hex characters (`GET {hub}/orgs/{orgId}/keys/{projectId} → {"projectId": "...", "key": "<hex>"}`). Once the project's key has been [rotated](#), the response also carries every generation: `{"projectId": "...", "key": "<generation-1 hex>", "currentGeneration": 2, "keys": [{"generation": 1, "key": "<hex>"}, {"generation": 2, "key": "<hex>"}]}`. `key` is always generation 1. The desktop engine may cache the key locally, protected by DPAPI on Windows or file permissions on macOS/Linux. With the keys and the stored blocks, you can decrypt the project yourself.
- **E2E:** the client creating the project generates the key and wraps it separately to each enrolled device's X25519 public key, and to the org's recovery public key if one is configured. Each wrapped copy is an envelope: `[0x01][nonce:12][ephemeral X25519 public key:32][AES-256-GCM ciphertext + tag]`, where the AES key is HKDF-SHA256(X25519(device private key, ephemeral public key), salt = nonce, info = `"swarmfile-e2e-envelope/v1"`). The org recovery private key is split into Shamir shares that are handed out when it's created and never stored by the hub. A quorum of shares recovers the project key. See [End-to-End Encryption](#).

Key rotation

A managed project's key can be **rotated** by the project's owners and admins (`POST /keys/:projectId/rotate`). Rotation creates a new key **generation** and makes it current. It never replaces or deletes an earlier generation:

- Every block records the generation that sealed it in its **version byte**. Everything written before a project's first rotation is generation 1 (`0x01` or `0x02`), so a rotation doesn't change or re-encrypt any stored block.
- New blocks are sealed under the current generation (`0x80 + g`). Readers look up the generation the byte names and use that key. There's no trial decryption.
- Every retained generation is stored wrapped under the hub's KEK, like the generation-1 key. A **KEK rotation** re-wraps all of them.
- Retired generations are kept so sealed blocks stay readable; deleting one makes its blocks unreadable.

E2E projects rotate the same way. The generation model and the version byte are identical (`0x80 + g`), and retained generations stay readable, so nothing here is managed-specific. The difference is *who holds the keys*: the hub never sees an E2E key, so a member's own device performs the rotation — it mints a new generation and re-wraps it to every remaining device and to the org recovery key — and `POST /projects/:id/rotate-key` only stores the envelopes the device produced. A device that doesn't yet hold the new generation fetches it on its next check and keeps its older-generation key for existing content.

A reader that predates rotation rejects a `0x82 +` block as an unsupported version. It never decrypts one into wrong bytes. Such a reader still opens every generation-1 block, because `key` in the key response stays generation 1. For the same reason, an older engine keeps writing `0x02` blocks under generation 1, which every reader can open.

Why a marker in the version byte, and not trial decryption. Trying each generation's key until the GCM tag verifies would also have worked without a format change. But AES-GCM checks its tag only after processing the whole block, so every read of an older block would cost one full decryption per newer generation, and a missing key would look like a corrupt block. Putting the generation in the byte that already selects the key derivation keeps the layout and the 29-byte overhead unchanged, costs nothing on reads, and makes a missing key a specific, recoverable error. The trade-off is a limit of 64 generations per project.

Hub-generated preview images

On a `managed` project, the thumbnails, video poster frames and point-cloud preview images the hub generates are JPEGs stored with the **block ciphertext layout** and the project key. Where a content block uses its CID, a preview uses the **SHA-256 of the JPEG bytes**, as 64 lowercase hex characters. That id is both the HKDF salt and the object name, at `p/{projectId}/{sha256}` (behind the same org prefix or bucket as **blocks**). To read one, decrypt it like a content block (a `0x02` block, or a rotated generation's `0x82 +` block) with the SHA-256 hex standing in for the CID, then check that `sha256(plaintext)` equals that hex.

A preview written before this scheme is a plain JPEG at the bare key `{sha256}`. Readers try the scoped key first, then the bare key. On the `plaintext` tier, previews are always plain JPEGs at the bare key.

Erasure coding

Some chunks also get **Reed-Solomon** erasure coding over $GF(2^8)$, with **10 data shards + 4 parity shards**. Any 10 of the 14 shards rebuild the chunk.

When it's applied. Erasure coding is *not* applied to every block. The engine decides for each save (`SWARMFILE_EC_ENABLED = on, off, or auto`, default `auto`). In `auto` mode, shards are produced only for files of **at least 64 KiB**, and only when the machine has a peer-to-peer fabric with other peers, favoring WAN and seed peers. It's skipped when two or more LAN peers are present, unless deliberate LAN shard placement is turned on, and when WAN peers are measured as fast (p50 RTT $\leq$ 80 ms). Uploads with no P2P fabric (cloud-only mode, and bulk migration by default) don't produce shards. These rules may change. A reader should just follow each `ChunkRef`'s `erasure` field.

Shards are extra. The whole chunk block is always stored too. Shards add redundancy and are never a substitute for the chunk.

Encoding. The input is the chunk's **stored** bytes: the ciphertext envelope on an encrypted project, or the plaintext otherwise.

1. `shard_size = ceil(len(stored) / 10)`.
2. Zero-pad `stored` to `10 × shard_size` and split it into data shards 0–9.
3. Compute parity shards 10–13 with Reed-Solomon (10, 4) over $GF(2^8)$. The engine uses the Rust `reed-solomon-erasure` crate (`galois_8`). Its exact encoding matrix is defined by that crate and isn't restated here.
4. Each shard is stored as a block whose CID is `blake3(shard bytes)`. Shards are **not** encrypted again.

erasure object (on a `ChunkRef`):

```
{
  "original_cid": "9f1c...e2a0",
  "original_size": 1210363,
  "shard_size": 121037,
  "shards": [
    {"cid": "...", "shard_index": 0, "shard_size": 121037},
    ...
    {"cid": "...", "shard_index": 13, "shard_size": 121037}
  ],
  "encrypted": true
}
```

Field	Meaning
<code>original_cid</code>	The chunk's CID, the same as the <code>ChunkRef</code> 's <code>cid</code> .
<code>original_size</code>	Length of the stored bytes that were encoded, so on an encrypted project it includes the 29-byte envelope. It is not the plaintext size.
<code>shard_size</code>	Shard length, the same for all 14.
<code>shards</code>	Exactly 14 entries. <code>shard_index</code> 0–9 are data shards and 10–13 are parity.
<code>encrypted</code>	Whether the encoded bytes were ciphertext. If absent, it defaults to <code>false</code> .

Reconstructing from shards: collect at least 10 of the 14 shards, checking each against its CID. Run Reed-Solomon reconstruction, join data shards 0–9 in order, and truncate to `original_size`. The result is the chunk's stored bytes. Decrypt and verify them exactly as you would a directly fetched chunk block.

Where blocks are stored

Blocks sit in an S3-compatible bucket, and the object key is built from the CID. The layout depends on the org's storage tier:

Storage tier	Object key for a project's block
Shared bucket (Free plan and default)	<code>{orgId}/p/{projectId}/{cid}</code>
Dedicated bucket (org-exclusive bucket)	<code>p/{projectId}/{cid}</code>
Bring your own storage	<code>{configured prefix}/p/{projectId}/{cid}</code> , or <code>p/{projectId}/{cid}</code> with no prefix configured

This applies the same way to chunks, manifests, segments, roots, and erasure shards.

Bare keys. Some blocks — written by clients that didn't name a project — sit at the bare key `{cid}` (behind the same org prefix or bucket as above). Readers try the project-scoped key first and fall back to the bare key. Take the **scoped key first**: on encrypted tiers, a bare object with the same CID can hold another project's ciphertext, which this project's key can't decrypt.

Other keyspaces, under the same org prefix or bucket:

- `lfs/{projectId}/{oid}`: directly uploaded git-LFS objects, addressed by git-lfs's `oid` (the SHA-256 of the file, 64 hex characters). Each is **one plaintext object** holding the whole file. A compatible older layout may also exist: the object at that key has custom metadata `enc: managed-v1` and holds a small JSON `{"enc":"managed-v1","chunks":N,"size":S}`, with 4 MiB ciphertext pieces at `lfs/{projectId}/{oid}/c/{n}`. Each piece uses the block ciphertext layout above, with the ASCII string `{oid}-{n}` (`n` counting from 0) taking the place of the CID in the key derivation.
- `t/{orgId}/{projectId}/{cid}` and `c/{orgId}/{projectId}/{hash}`: tree and commit objects for version history. Their format isn't covered here.
- `p/{projectId}/{sha256}`: [encrypted preview images](#) for a managed project. They share the project-scoped namespace with blocks but are named by SHA-256, not BLAKE3.
- `{sha256}` (bare), `_videoproxy/`, `_materialized/`: unencrypted derived renderings (see [above](#)).

The physical bucket name, the account behind it, and anything else about the shared bucket's arrangement are operational details. They aren't specified here and may change.

Reconstructing a file by hand

You need:

- the file's `rootCid`, from the project's metadata;
- read access to the stored objects (your dedicated or BYOS bucket, or blocks fetched through the hub's block API);
- for a `managed` or `e2e` project, the 32-byte project key, or every retained generation of it if the project's key has been [rotated](#).

Then, in this order:

- Define `fetch(cid)`**: read the object at the [project-scoped key](#), falling back to the bare key. Then:
 - no project key (plaintext tier): the stored bytes are the plaintext;
 - with a **managed** key: if `blake3(stored) == cid`, the block is plaintext (a tolerance for managed projects only — an E2E or static-key reader must treat a plaintext block as an error), so use it as-is. Otherwise its first byte must be `0x01`, `0x02`, or `0x82 – 0xC0`. Pick the key generation that byte names (generation 1 for `0x01 / 0x02`, byte `– 0x80` otherwise), derive the per-block key for that version, and AES-256-GCM-decrypt `stored[13..]` with nonce `stored[1..13]` (the tag is the last 16 bytes).
 - Verify** that `blake3(plaintext) == cid`. Stop on any mismatch or GCM authentication failure.
- Fetch the root**: `root = JSON.parse(fetch(rootCid))`.
- Classify it** using the rules in [Telling the shapes apart](#).
 - Flat: `chunks = root.chunks` (or a `root.blocks` array).
 - Segmented (`manifest_version == 2`): for each entry in `root.segments`, in order, run `seg = JSON.parse(fetch(entry.cid))`, check that `len(seg.chunks) == entry.chunk_count`, and append `seg.chunks`.
- Fetch each chunk**: for each `ChunkRef` in order, `data = fetch(ref.cid)` and check that `len(data) == ref.size`. If a chunk block is missing and the ref has an erasure group, rebuild the chunk's stored bytes from any 10 shards (see [above](#)), then decrypt and verify them as in step 1. A block 404 is not proof of loss: an in-flight save's manifest can reference blocks that have not been uploaded yet, so re-fetch after the save completes before treating a block as permanently missing.
- Assemble**: write each chunk's bytes at `ref.offset`. Since chunks are contiguous and in order, you can simply concatenate them.
- Check the total**: the output length must equal `total_size`. The manifest has no separate whole-file hash. Integrity comes from the per-block CID checks above, and the manifest itself is covered by `rootCid`.

A 1-byte range read only needs the chunks that overlap it. With a segmented root, binary-search `first_offset` to find which segments to fetch.

Versioning & stability

Stable: a stored object is never rewritten to upgrade its format, so existing data stays readable without migration. Changing anything below would change CIDs for existing content, so the code treats each of these as a breaking change:

- BLAKE3-256 hex CIDs over plaintext;
- the flat manifest's field names and serialization;
- the segmented root and segment shapes, with `manifest_version: 2` ;
- the `[version] [nonce] [ciphertext+tag]` envelope and the `0x02` key derivation;
- the key-generation version bytes `0x82 – 0xC0` , which use the `0x02` derivation under that generation's key;
- readers accepting the older manifest shape and `0x01` ciphertext.

How format versions are signalled:

What	Signal
Flat vs segmented manifest	Which fields are present, plus <code>manifest_version</code> (absent = v1, <code>2</code> = segmented). Readers reject unknown versions.
Ciphertext envelope	The first byte (<code>0x01</code> version 1, <code>0x02</code> current, <code>0x80 + g</code> for key generation $g \geq 2$). Readers reject unknown versions.
Chunking	The manifest's <code>chunking.type</code> (<code>fixed</code> / <code>cdc</code>).
E2E key envelope	Its first byte (<code>0x01</code>).

May change without notice (none of this affects reading data that's already stored):

- the default chunking parameters and the per-extension profiles;
- when erasure coding is applied;
- the block size limit and the $\frac{3}{4}$ segment target (which change which files get segmented and what their manifest CIDs are);
- bucket and prefix arrangement beyond the key shapes documented above;
- how git-LFS objects are stored.

Not specified here: file and folder metadata, tree and commit objects, and anything about Reed-Solomon beyond "10+4 over GF(2⁸) via `reed-solomon-erasure`". If you depend on one of these, [contact us](#).

Telemetry & Data Collection

What Swarmfile measures, what leaves your machines, and what it deliberately does not collect. This page is for security reviews and procurement questionnaires; the network-level view of every connection is in [Network Requirements](#).

Short version

- **No third-party analytics or crash-reporting SDKs ship in the desktop app or engine.** There is no telemetry sent to an analytics service, ad network, or crash collector from the client. The **website is separate**: Google Analytics is scoped to the public marketing/docs pages and is **not loaded on public project pages, Explore, the signed-in dashboard**, or any app route — so no project slug, file path, or dashboard URL is reported to it. It also records simple interaction events such as download-button clicks (which file, which page — no account id). (Navigating from a marketing page into any of those becomes a full page load, whose document has no analytics at all.) The Google Fonts stylesheet, by contrast, loads site-wide (the dashboard uses the same type stack); its requests necessarily carry the page URL but nothing else.
- **OpenTelemetry only exists in Swarmfile-operated diagnostic builds, and only if you point it somewhere.** Shipped desktop installs don't emit traces. In a diagnostic build, an unset endpoint still defaults to a local OTLP collector at `http://localhost:4318` ; set it explicitly if you want no export attempts at all.
- **File contents never leave unencrypted** on a managed or end-to-end project — the storage layer and peer transfers see ciphertext, with the documented exceptions on [Security Architecture](#).
- **Usage metering counts bytes and seats, not contents.** Billing dimensions are computed from byte counters and org membership.

Client-side diagnostics (opt-in)

Mechanism	Default	What it sends	Where it goes
<code>SWARMFILE_OTEL_ENDPOINT</code>	<code>http://localhost:4318</code> in Swarmfile-operated diagnostic builds; unset on shipped installs, which don't emit traces	Traces/spans for engine operations: operation names, timings, counts, error kinds, and internal identifiers (entry ids, content hashes). No file bytes. Diagnostic builds only : shipped desktop installs don't emit traces.	The OTLP collector you configure
<code>SWARMFILE_OTEL_SERVICE_NAME</code>	<code>swarmfile-engine</code>	Service name on those spans	Same collector
<code>SWARMFILE_UPLOAD_TRACE</code>	Off	Promotes upload-attempt logging (including the hub's refusal body) to warn in the local log	Local log file only
<code>swarmfile-doctor</code> / <code>swarmfile doctor</code>	Runs when you ask	Probe results, hostnames, versions, and error text, written to a local JSON report	Stays local until you attach it to a support request
<code>RUST_LOG</code>	Off (info)	More detailed local logging	Local log file only

Normally the client also sends small **operational reports** to the hub — a P2P bandwidth delta roughly every 60 seconds, a branch-switch report when a machine changes branch, and the metadata calls that make up ordinary work. These carry counters, byte totals, and ids; never file content.

What the service collects

Usage metering (billing and limits)

The hub meters usage per org:

Dimension	Measured	Billed
Storage	Bytes stored per project (from committed manifests and object sizes)	Overage, then hard ceiling
Egress	Bytes served, counted per direction (cloud egress and P2P in/out are tracked separately)	Not billed — unlimited on every plan; a separate platform-wide hourly ceiling can refuse further anonymous raw reads under extreme volume (it isn't plan- or org-specific and can't be raised)
Collaborators	External collaborators on private projects above the per-seat allowance (public-project guests are free and uncounted)	Overage, then hard ceiling
Seats	Members counted per org and reconciled to the subscription	Per seat
Headless API keys	Keys minted, by org and project	Hard cap; key packs raise it

Metering is counters and sizes — it never inspects file contents. Details, allowances, and ceilings are in [Billing & Plans](#).

Installer downloads and installs

- **Anonymous downloads.** Installer links on public pages (the download page, the welcome page, the get-started banner, and the Security page's doctor link) go through a hub redirect that records one event per fetch: which file, which page/button, and when. The event stores no token, org, user, or device identifier; the per-IP rate limiter keeps only a short-lived counter to bound abuse. Rows are deleted after 180 days.
- **Signed-in downloads.** A download from the dashboard's Setup tab goes through a per-seat token and records which org and seat fetched which installer and when (also deleted after 180 days), so a download can be tied to the seat that made it.
- **First-launch install report.** On first launch the Desktop App reports the machine id, OS, and app version, and refreshes a last-seen time on later launches, so the Setup tab can show which of your machines are connected. The server may link that report to a recent installer download from the same account; nothing is added to the installer itself.

These records never include file contents. Aggregate download counts are available to support for capacity planning.

Activity and audit events

The org **Activity** feed (and its CSV export) records user-visible metadata per event: `occurredAt` (the CSV column is `occurred_at_iso`), `kind`, `actor` (user id, display name, email), `project` (id, name), `entry` (id, name), a summary, and a detail field. Examples are commits, merges, permission changes, quarantine decisions, and plan/policy changes. It does **not** record file contents, and it never records a file's bytes.

Retention applies to these rows — see [Audit-event retention](#). The personal notification inbox and its email copies contain the same class of metadata (who, what, when), not content.

Notifications by email

Email delivery uses **Postmark** as the transport provider. The message contains the same event metadata as the in-app notification (project and entry names, actor names, a link), never file contents. Providers and hosts are listed in [Network Requirements](#).

Third parties that can see what

Provider	What it handles	What it can see
Cloudflare (Workers, R2, Pages)	Hub, identity, storage hosts	Encrypted blocks for private projects; plaintext for Free-plan/public projects (by design); request metadata
Postmark	Notification and account email transport	Recipient address and the email's metadata content
Stripe	Billing	Org/account details and payment method, never file content
iroh/NO relay network	NAT traversal for P2P when direct connections fail	Connection metadata (endpoints, timings); relayed traffic is encrypted end-to-end between peers
Google (swarmfile.com website)	Google Analytics on the marketing/docs pages; Google Fonts site-wide	Page views, the referring URL, and download-button click events (which file and which page); nothing from the app or your projects

What is not collected

- No browsing history in the dashboard, no cursor tracking, no session recording.
- No content indexing of your files by the service beyond what a preview/search feature needs (managed tier), and none at all on end-to-end projects.
- No telemetry from the Desktop App UI itself.
- No automatic upload of logs or doctor reports — support only gets what you attach.

Verifying this on a machine

1. Confirm no OTLP traffic leaves the host: shipped desktop installs don't emit traces, so ports 4317/4318 stay silent even if the endpoint variable is set.
2. Run `swarmfile-doctor --json --report /tmp/report.json` and inspect the file: it contains hostnames, versions, probe verdicts, and error strings.
3. Inspect the cache directory's logs for what `RUST_LOG / SWARMFILE_UPLOAD_TRACE` would add.
4. For an end-to-end project, confirm the mount's blocks are ciphertext by reconstructing one as described in [Storage Format](#).

Where to go next

- [Security Architecture](#) — encryption tiers and threat model.
- [Operations](#) — the audit log and its CSV export.
- [Trust & Compliance](#) — DPAs, subprocessors, and incident response.

Known Limitations

Swarmfile covers a wide surface — a mounted drive, version control, git, LFS, identity, billing — and some of it is deliberately not built, or works differently from the tool it replaces. This page collects the limits that most often decide a fit, stated plainly rather than left for you to discover during a rollout. Each entry links to the page that owns the detail.

If a hard limit here is a blocker for your team, [tell us](#) — it is how we decide what to build next. For the exhaustive per-subsystem gap list, see [Filesystem Compatibility → Declared gaps](#).

What revoking access can and can't reach

Revoking access — a removed grant, a removed member, a revoked guest, or a revoked device key — is enforced by the hub on every request immediately, and affected machines are told to delete their cached copy of the revoked files (removed from the drive's listing, cached file data purged, and a purge receipt the organization can see under **Settings → Access revocations**). Three boundaries are worth knowing before you plan a rollout around it:

- **A machine that never reaches Swarmfile can't be told.** An offline computer deletes its cached copy the next time it reaches the hub. The organization bounds that delay with the [offline access window](#), which is **on by default (7 days)**: a machine that hasn't reached the hub within the window pauses local access (even to kept-offline files) until it reconnects. An organization that turns the window off has no upper bound on how long a disconnected machine keeps serving what it already cached.
- **Deletion is not cryptographic erasure on the managed tier.** Cached files on managed and plaintext projects are deleted; on end-to-end encrypted projects the key is dropped so what remains is unreadable. In neither tier can Swarmfile reach bytes an application already read (they live in that app's memory), the operating system's page cache, or copies made by a third-party previewer or backup tool.
- **A removed member's own unsynced saves are kept.** Work they saved that never uploaded stays on their computer, held rather than deleted, and uploads if access is restored (on encrypted tiers it is unreadable once the project key is dropped). Removal revokes access; it does not destroy their unuploaded work.

No Unreal Engine or Unity source-control plugins

There's no Unreal Editor or Unity source-control provider for Swarmfile today, so those editors' built-in check-out and submit buttons won't talk to it. Teams work on the mounted drive directly and use the Desktop App or the `swarmfile` CLI for locks and commits.

This is separate from Revit/SolidWorks-class worksharing, which needs no plugin on Windows: there, Swarmfile honors the native OS exclusive-open lock the application already takes and turns it into a hub-enforced claim across machines. Detail: [Moving from Perforce → Not supported](#).

No Perforce history importer

There is no tool that reads a Perforce server's revision history and replays it as Swarmfile commits. A migration brings the **head revision** across; history in Swarmfile starts from there, and the old Perforce server stays around read-only for anything older. If you need the old revisions auditable, plan to export or retain them rather than decommissioning the server on cut-over. Detail: [Moving from Perforce → Not supported](#), [Migrating Existing Data](#).

No git rebase or history rewriting

Swarmfile history is an append-only, server-side DAG of whole-file commits. It isn't rewritten — you can't squash, rebase, or clean up a run of autosaves into one tidy commit.

What you can do instead: `swarmfile commit --amend -m "..."` changes the most recent commit's **message only** (contents and history are untouched), and `describe` or `tag` puts a name on any past commit, including a bare autosave. `restore` and `revert` move history *forward* as new commits rather than erasing it. Detail: [Coming from Git → What genuinely has no equivalent](#), [Git and Swarmfile → Compatibility at a glance](#).

Importing an existing repository converts or skips some git shapes

`swarmfile git import` brings a GitHub, GitLab or local repository into a project, but the git view is not a full git server, so a few things are converted rather than carried byte-for-byte:

- **Annotated tags** have no equivalent and are imported as lightweight tags on the same commit (the default; `--annotated-tags skip` leaves them out, `fail` stops before creating anything) — tag messages and signatures are not preserved.
- **A branch the git view cannot hold** — one containing submodules (`.gitmodules`), git-LFS pointers, or an octopus merge — is skipped and listed by the import; if the problem is on the source's *default* branch, the import refuses before creating the project.
- **Upstream rewriting and deletion stay out of sync.** Re-running the import adds new commits, branches and tags; a force-pushed (rewritten) upstream branch is refused rather than mirrored, and branches or tags deleted upstream are not deleted in the project — remove those there with the normal branch and tag commands.
- **Git-LFS objects aren't fetched by the importer.** An LFS-using default branch is refused and an LFS-using side branch is skipped, rather than importing pointer files with no objects behind them.

What you can do instead: keep the source as history-of-record until the import is verified, and prefer `git push / git-LFS` as the ongoing bridge when the repository keeps moving. Detail: [swarmfile git import](#), [Git and Swarmfile → Compatibility at a glance](#).

git-LFS isn't available on end-to-end encrypted projects

A stock `git-lfs` client can't hold the per-user key, so git-LFS is refused on end-to-end encrypted projects. It works on managed-key and unencrypted projects, with one carve-out worth knowing: objects uploaded directly by the stock `git-lfs` client — and uploads of 64 MiB or more through the Desktop App's built-in LFS agent — are stored **without** application-layer encryption. Smaller uploads through the built-in agent, and every upload through the standalone `swarmfile-lfs` agent, are

encrypted like any other block. If a binary must be encrypted at rest, keep it on the mounted drive rather than routing it through LFS. Detail: [Git-LFS → Which projects it works on](#), [Security → Encryption](#).

Git materialization follows your access today

A clone or fetch reads the whole of every commit it walks — each commit's tree exactly as it was archived — and access is evaluated against the access you hold **today**:

- A folder restricted for you **now** means any commit whose tree contains it can't be handed over whole; a clone or fetch that has to walk one is refused with a message saying so, rather than returning a repository with silent holes. Grant read and the same clone succeeds — including commits made while you were restricted; materialization is not frozen to the access of the commit's date.
- Restricting a folder retracts it from clones at once, for new commits and old ones alike.

The mounted drive, [swarmfile materialize](#) and git all follow the access you hold today. The access that existed when a commit was made is still recorded, and an explicit point-in-time history view can ask for that replay — it is not what git reads. Detail: [Clone a project → What a clone contains](#).

A release copy is a snapshot of files and directories

Copying a release — the web **Copy to my org** action or `swarmfile clone ... --into` — lands the release's files and directories as a new snapshot on the destination's **default branch**. It is not a fork of the project: no history, no tags, and no later sync from the original.

- The web action always creates a **public** project (it copies a public release). For a **private** destination, copy through the Desktop App's CLI instead — it seals the blocks on your machine before they upload (see [Publishing releases → Getting a copy](#)).
- Symlinks are not part of a public release, so a copy contains none; directories are recreated from the manifest's paths.
- A copy adds files; it never overwrites. A same-name file with different content refuses the copy, and a destination whose default branch is protected refuses it (409) until the protection is lifted.

Encryption tier is fixed at project creation

Managed (the default), end-to-end, and plaintext (public) are chosen when a project is created and never change afterward. Upgrading a plan, or deciding later that you want E2E, does not retroactively encrypt an existing project — you create a new project at the tier you need and bring the content across. Choose deliberately at creation if encryption posture matters. Detail: [Getting Started → Creating a project](#), [Security Architecture → Unencrypted tier](#).

SSO sign-in needs your org's slug

There's no email-domain auto-detection. A member signing in with SSO enters their **organization's slug** (the short name in the org's Swarmfile URL), or follows a direct per-org sign-in link an admin can share — rather than typing an email and being routed to the right identity provider automatically. It is a one-step onboarding cost for non-technical staff; the per-org link removes it on a machine an admin provisions. Detail: [Getting Started → Signing in, Identity](#).

Switching accounts requires an app restart

Switching to a different account while the app is running is refused (`account_switch_requires_restart`). Any unsent work from the previous account stays on your machine, quarantined rather than lost. The requested sign-in is kept: restart when the Desktop App offers it and the switch completes. This is most visible to freelancers and contractors working across several client orgs in a day. Detail: [Getting Started → Signing in](#).

swarmfile

`swarmfile` drives Swarmfile version-control operations against a running engine: commit, switch/restore, branches, merge, and the rest of the day-to-day project commands. It talks to the engine over its control socket — every subcommand mirrors an operation the Desktop App can do, and every subcommand requires a running engine except `completions`, `demo offline` (it changes the host firewall, not the engine, and must work exactly when the engine is what is being cut off), the `git` helper verbs (`git url`, `git clone` and `git import`, which answer before the engine connection is needed), and the pure Git-habit shortcuts (`push` / `pull` / `clone` / `rebase` / `cherry-pick` / `rm` / `mv` / `cp` / `mkdir` / `touch` / `worktree` / `grep` / `init` / `remote` / `clean` / `reflog` / `mergetool` / `difftool` / `shortlog` / `submodule` / `bisect` / `gc` / `prune` / `fsck` / `archive` / `not` which work standalone. `add`, `stash`, `reset`, and the executable `clone <id> <path>` / `clone ... --checkout` / `clone ... --into` forms are the exception among the shortcuts — they map onto real commands, so they need a running engine too (see the Git-habit shortcuts note below).

If the engine isn't running, start it (or the Desktop App) first, or use the standalone `swarmfile-doctor` binary to diagnose connectivity without one.

`swarmfile` comes with the desktop app on every platform: `/usr/bin/swarmfile` from the Linux `.deb`, `/usr/local/bin/swarmfile` on macOS (a symlink into `Swarmfile.app`, so desktop updates carry it too), and `swarmfile.exe` in the Windows install folder (`C:\Program Files\Swarmfile`), which the installer adds to the machine `PATH`. There is no separate CLI download — install the desktop app and open a new terminal. An install from before the tools shipped (or a macOS install that came from the `.dmg` rather than the `.pkg`) needs one full installer run to get them and the `PATH` entry — **Diagnostics → Reinstall** does that from the tray; on macOS the tray also repairs the `/usr/local/bin` symlinks at launch where the directory is writable.

`swarmfile` doesn't include a `search` subcommand — `search` needs its own hub-authenticated round trip independent of the control socket, so it lives in the standalone `swarmfile-search` binary instead.

Global flags

These apply to every subcommand.

Flag	Description
<code>--cache-dir <path></code>	Engine cache dir (where <code>control.json</code> lives). Respects <code>\$SWARMFILE_CACHE_DIR</code> . Default <code>~/.cache/swarmfile</code> (macOS/Linux), <code>%LOCALAPPDATA%\swarmfile</code> (Windows).
<code>--json</code>	Print raw JSON instead of a human-readable summary.
<code>--mount <id></code>	Target this mount instead of whichever one is <code>current</code> — see <code>mounts</code> below. Has no effect on <code>mounts</code> subcommands themselves, which address by their own <code>id</code> argument instead, nor on <code>workspace switch</code> , which changes the whole engine's active workspace rather than one mount.
<code>--no-wait</code>	Don't wait for a <code>long operation</code> : print its id and exit <code>3</code> at once.

These are global flags, so they work on either side of the subcommand: `swarmfile --json status` and `swarmfile status --json` are equivalent, as are `swarmfile --mount m2 status` and `swarmfile status --mount m2`. That holds for the git-habit shims too — `swarmfile rm foo --json` is read as the flag, not a path.

Exit codes

Scriptable commands separate "you need to do something" from "something failed", so a CI or pre-flight script can branch on it:

Code	Meaning
0	Success — including the Git-habit shortcuts whose intent the product already satisfies (<code>push</code> , <code>pull</code>): nothing to do, so nothing failed.
1	Any other failure: an unreachable engine, a hub transport error, and the shortcuts whose operation has no equivalent here.
3	A <code>long operation</code> is still running — after <code>--no-wait</code> , or when you pressed Ctrl-C to detach. It carries on in the engine; follow it with <code>op status</code> / <code>op wait</code> .
2	A refusal a script can act on, or a CLI usage error (an unknown command, flag, or argument). With <code>--json</code> , refusals print their response body as before, and usage errors print <code>{"code": "usage_error", "kind": "...", "error": "..."} </code> instead of clap's prose — so a script can tell a typo apart from an actionable refusal.

A command page can assign `3` its own meaning where there is no long operation to follow — `git index` exits `3` when the hub couldn't be reached, for example.

The translated shortcuts (`add`, `stash`, `reset --hard <rev>`, `clone <id> <path>`) return the exit code of the command they map to, not a shortcut code.

Git-habit shortcuts honor `--json`: instead of the sentence they print `{"code": "already_satisfied"|"not_applicable", "exitCode": N, "hint": "..."}` .

The confirm-gated verbs all take `-y` / `--yes`, and most also accept `-f` / `--force` as a synonym — except `git enable`, and the verbs with a separate safety-gate `--force` below — `stash drop` / `clear`, `changelist discard`, `branch prune`, `trash purge`, `project delete`, `sync discard`, `git enable` (one-way, so it never proceeds without it), `ignore-clean`, and (unless `--dry-run`) `cache reclaim` / `scratch gc`. Without it, a non-interactive stdin or `--json` refuses (`{"code": "confirmation_required"}`, exit 2) rather than destroying work; the estimate-only verbs (`ignore-clean`, `cache reclaim`, `scratch gc`) take `--dry-run` to preview without confirming or deleting. Two prompt-gated verbs are the exception and proceed without `--yes` in the right context: `hydrate start` and `workspace switch` both proceed under `--json` or a non-interactive stdin (an interactive terminal still prompts). Both are resumable/reversible operations rather than destructive ones, and `--yes` only skips the interactive prompt.

A few verbs have a separate `-f` / `--force` meaning "proceed despite a safety gate" — `mounts close` (open changelists), `scratch clear` (project still open), and `ignore-clean` (`-y` skips its prompt, `-f` also overrides its recent-foreign-edit gate).

Commands that currently use exit 2: `mr status` (open, not yet approved — the CI review gate), `conflicts resolve --auto` / `--all --auto` (a human must choose — a request that *fails* mid-sweep is exit 1 and takes precedence), `lock` (already held), `mounts close` (open changelists, or the boot mount), `acl mode set` / `acl entry ...` (`forbidden` / `not_found` / `plan_restricted`), `unlock-request respond` / `cancel` (already resolved, or not yours), `workspace switch` (another switch running, a changelist open, or a drive open on another project — `cross_project_mount_open`), `checkout` with no target, `ignore-clean --yes` when a candidate was recently edited by someone else (unless `--force`) or when any delete failed, and `hydrate start` / `fetch` when a job is already running, and any confirm-gated verb without `--yes` (the `confirmation_required` refusal above; `ignore-clean --yes` keeps its own cross-machine gate).

Version control

See [Version Control](#) and [Branches & Merging](#) for the conceptual walkthrough — this section is the flag-by-flag reference. If you're arriving from Git, [Coming from Git](#) maps your habits onto these commands.

Git-habit shortcuts. Typing a Git command that has no Swarmfile equivalent — `push`, `pull`, `clone`, `rebase`, `cherry-pick`, `rm`, `mv`, `cp`, `mkdir`, `touch`, `worktree`, `grep`, `init`, `remote`, `clean`, `reflog`, `mergetool`, `difftool`, `shortlog`, `submodule`, `bisect`, `gc`, `prune`, `fsck`, `archive`, `notes` — prints a one-line pointer to the real workflow instead of an error (e.g. `push` explains that saves sync automatically, `rm / mv` point at doing it on the mounted drive, `worktree` points at `swarmfile mounts open`, `grep` points at `swarmfile-search`, `mergetool` points at `conflicts resolve --tool`, and `clone` is argument-aware — a public-project URL, a `/s/<token>` share link, a `swarmfile://mount?project=...` deep link, or a lone project id each get the right next step; `clone <project-id> <path>` actually opens a mount — add `-b <branch>` to pin it to that branch, exactly like `mounts open --branch`, or add `--checkout` to write a plain directory instead (see `materialize`); `clone <public-project> --into <org>/<project>` copies a public release into a project, encrypting blocks on this machine — the way to get a public release into a **private** project, since the web "Copy to my org" fork only produces public ones. Add `--tag <tag>` to pick a release, `--path <p>` (repeatable) to copy a subtree, and `--new` (with optional `--public`) to create the destination project first. It runs as a long operation (`--no-wait`, `swarmfile op status|wait|cancel`; per-file progress; a cancel resumes from what's done), and its refusals (unknown slug or tag, no write access, a protected branch, a name collision) exit `2` while hub or transport failures exit `1`. These shortcuts work even when the engine isn't running (the executable `cclone` forms need it).

Three shortcuts have a real destination and **run it** instead of just advising: `stash` (and `push / save`) → `changelist park`, `stash list / show` → `changelist parked`, `stash pop / apply` → `changelist resume`, `stash drop / clear` → `changelist discard` (`stash list` numbers entries `stash@{n}`; `stash drop stash@{n}` or `drop n` discards just that one); `add` → opens a changelist if none is open (idempotent — saves then stage into it automatically, and this mount goes into staged mode until you `commit` or `changelist cancel`); and `reset --hard <rev>` → `restore <rev>` (one undoable commit). `add <path>` exits `1` if the path doesn't resolve, so a typo doesn't silently pass. Parked work is a **set, not a stack**: a bare `pop / apply` resumes all parked work for the branch, and `pop <n>` (or `pop stash@{n}`) resumes just that entry and leaves the rest parked — an entry parked on another branch is refused, since parked work resumes on its own branch (switch to the new id). If an entry's changeset was closed on the hub while it sat parked, resuming re-creates it and its staged files under a new id, and the result names the new id. A bare `drop` or `clear` abandons all of it — so discards ask first (or take `--yes`). On success these print a sentence, not raw JSON. Bare `reset --hard`, `--soft / --mixed`, and bare `reset` still print guidance, since there is no unstaged state to discard. These three need a running engine.

Referring to a commit. Every command that takes a commit (`show`, `restore`, `revert`, `describe`, `diff`, and the commit form of `checkout`) accepts any of: a commit number (`482`), the `#N` form the log prints (`#482`), `HEAD`, `HEAD~N` / `HEAD^` (`N` commits back in the current branch's log), a **tag name**, a **branch name** (its head), a changeset id, or a 64-hex **commit hash**. A `seq: / head: / branch: / tag: / changeset: / hash:` prefix forces which kind is meant. Resolved client-side to the commit number before the request.

commit

Open a changelist (if none is open) and submit it in one step.

Flag	Description
<code>-m</code> , <code>--message <String></code>	Commit message
<code>--amend</code>	Rename the most recent commit instead of creating a new one (like <code>git commit --amend</code>). Message-only and non-destructive — contents and history are untouched. Requires <code>-m</code> . Equivalent to <code>describe</code> on the latest commit.

```
swarmfile commit -m "Fix layout regression in the deck"
swarmfile commit --amend -m "Fix layout regression in the title deck"
```

changelist

Explicit changelist control (Mode B: staged atomic submit) — for when you want to stage multiple saves before committing them atomically, rather than using `commit`'s one-step flow. A mount can hold several open changelists at once (Perforce's `-c` model) — see [Multiple changelists at once](#) for the conceptual walkthrough.

Subcommand	Flags	Description
<code>open</code>	<code>-m</code> , <code>--message <String></code>	Open a new changelist and make it current; subsequent saves stage into it until submit/cancel/switch. Never fails if one is already open — opening a second doesn't close the first.
<code>submit</code>	<code>--id <changeset></code>	Atomically apply a changelist. Defaults to the current one; <code>--id</code> targets a specific one without switching to it first.
<code>cancel</code>	<code>--id <changeset></code>	Abandon a changelist. Defaults to the current one; <code>--id</code> targets a specific one without switching to it first.
<code>status</code>	—	Show whether the current changelist is open and what's staged, plus any files blocked by an unresolved conflict — the same changelist + conflicts pair <code>status --changes</code> renders.
<code>list</code>	—	List every open changelist, <code>*</code> marking which one is current.
<code>current</code>	<code><id></code>	Switch which open changelist is current.
<code>reopen</code>	<code>[path]</code> , <code>--to <changeset></code> , <code>--from <changeset></code>	Move staged file(s) to a different open changelist without unstaging/restaging. With <code>path</code> , moves that one file. Without it, moves every staged file — from the current changelist, or from <code>--from <changeset></code> if given — to <code>--to</code> . <code>--from</code> applies only without a path (a path with <code>--from</code> is a usage error).
<code>park</code>	—	Set aside every open changelist on this mount (kept, and kept alive on the hub) so the branch/project can be left without abandoning the work.
<code>parked</code>	—	List parked changelists with age + files.
<code>resume</code>	—	Restore parked work for the mount's current branch.
<code>discard</code>	<code>--changeset <id></code> , <code>--yes</code>	Abandon parked work (one with <code>--changeset</code> , else all). Asks first (it cannot be undone); <code>--yes</code> skips the prompt.

```

swarmfile changelist open -m "Reorganize asset folders"
# ...edit and save files...
swarmfile changelist status
swarmfile changelist submit

```

```

swarmfile changelist open -m "Swap all Act 2 plates to final grade"
swarmfile changelist open -m "Quick fix: wrong LUT on shot 12"
swarmfile changelist list
swarmfile changelist current a1b2c3d4
swarmfile changelist reopen "Renders/Act2/shot12.mov" --to e5f6a7b8
swarmfile changelist submit --id e5f6a7b8

```

switch

Switch which branch this mount works on (like `git switch`) — hot and in-process, no engine restart. Same operation as `branch switch`. "This mount" is the `--mount one`, else the current mount (`mounts current`): the branch checks, `-c` 's create, and the switch itself all act on that same mount. A mount that shares another mount's branch context can't switch on its own and is refused (`shared_branch_context`) — `-c` refuses before creating anything. The success line names the branch the engine reports it switched to.

Argument / Flag	Description
[name]	Branch to switch to. Omit to return to the project's default branch (usually <code>main</code>).
<code>-c</code> , <code>--create</code>	Create the branch first (forking from the current branch), then switch to it — like <code>git switch -c</code> . Errors if it already exists.

Refused while any changelist is open on that mount (exit 2 on the default mount; on any other mount the switch fails with the same reason, exit 1); submit or cancel all of them first.

```

swarmfile switch look-dev-b
swarmfile switch          # back to the default branch
swarmfile switch -c feature-x # create + switch

```

restore

Restore a scope (the whole project, a folder, an entry, or a mount-relative path) to a past commit — a true tree snapshot, landed as one undoable commit (like `git restore` , but whole-project by default). Locked files → exit 2, and a file the restore would change that has an unresolved conflict refuses with `conflict_unresolved` — resolve it first.

Argument / Flag	Description
[seq]	A commit to restore to — <code>#N</code> , <code>HEAD~N</code> , a tag, a changeset id, or a 64-hex hash (see "Referring to a commit"). A branch name restores this mount's branch to that branch's head tree (it does not switch branches — that's <code>switch</code>).
<code>--at-seq <N></code>	Commit number to restore to — instead of [seq] , not with it (both at once is a usage error).
<code>--path <path></code>	Restore only this mount-relative path.
<code>--folder-id <id></code>	Restore a specific folder by ID.
<code>--entry-id <id></code>	Restore a specific entry by ID.

```

swarmfile restore 482          # whole project back to commit #482
swarmfile restore 482 --path shots/seqA # just that folder

```

`--json` prints the engine's response plus `targetSeq` — the commit number the CLI resolved your ref to; the engine's own `seq` is the NEW commit the restore landed as.

checkout

A Git-habit shim mirroring `git checkout` 's overloading — prefer `switch` / `restore` , which say exactly what they do.

You type	What happens
<code>checkout <branch></code>	Switches to that branch (→ <code>switch</code>).
<code>checkout <N></code> (a number, <code>HEAD</code> , <code>HEAD~N</code>)	Restores to that commit (→ <code>restore</code>).
<code>checkout <tag> / <hash> / <changeset-id> / a prefixed ref (tag:x , seq:5 , hash:...)</code>	Restores to that commit — resolved exactly as <code>restore</code> resolves it. <code>branch:x</code> switches to <code>x</code> .
<code>checkout -b <name> [<target>]</code>	Creates a branch and switches to it — like <code>git checkout -b</code> . Forks from <code><target></code> (a branch, tag, or commit — same rules as <code>branch create --from</code>) or <code>--at-seq <N></code> , else from the current branch. <code>--path</code> / <code>--folder-id</code> / <code>--entry-id</code> don't combine with <code>-b</code> (a branch is whole-project).
<code>checkout --at-seq <N> [--path ...]</code>	Restores the version at that sequence number (the <code>--at-seq</code> form). Not combined with a <code><target></code> .

A bare number restores that commit **unless a branch of that name exists** — an existing branch wins even when its name is commit-shaped, as in git, and the CLI prints a note when it does. Force the branch with `switch <name>`, or force the commit with a `hash: / seq:` prefix.

show

Show one commit: its message, author, time, and the files it changed (like `git show`). Whole-file — Swarmfile versions binaries, so there's no line-level diff.

Argument	Description
<rev>	A commit — number, <code>#N</code> , <code>HEAD</code> , <code>HEAD~N</code> , or a tag (see "Referring to a commit" above).

```
swarmfile show 482
swarmfile show HEAD      # the latest commit
swarmfile show HEAD~2    # two commits back
swarmfile show v1.0      # a tag
```

show-tree

Fetch and print a tree object by its content-addressed cid (like `git ls-tree`) — one level, not recursive: each entry's kind (`d` directory, `f` file, `t` symlink), name, and content cid. Get a cid from `swarmfile show <rev>`'s `tree` line. A directory large enough to be stored segmented is listed a segment at a time — each segment is its own fetch, never one unbounded request for the whole directory. A **project member's** listing is filtered against the selected mount's branch view (the same ACLs a normal listing applies), so a cid taken from an old commit can come back partial or empty; owners and service callers see the object unfiltered.

Argument	Description
<cid>	A tree object's content-addressed cid (64 lowercase hex characters).

```
swarmfile show-tree a1b2c3...
```

revert

Undo a commit by landing a **new** commit that restores the prior versions of the files it changed (like `git revert`) — non-destructive, itself undoable, and it doesn't touch newer commits. Locked files → exit 2, and a file the revert would change that has an unresolved conflict refuses with `conflict_unresolved` — resolve it first.

Argument	Description
<seq>	The commit to undo — a number (<code>#N</code>), <code>HEAD~N</code> , a branch, a tag, a changeset id, or a 64-hex hash (see "Referring to a commit").

```
swarmfile revert 482
```

rollback

Restore **one file** to an earlier version, keeping the current version in history — the CLI form of the Desktop App's version picker. File-scoped and non-destructive, unlike `revert`, which undoes a whole commit.

Requires the file's write lock, and refuses (exit 2) if the file has uncommitted local changes, is locked, has an unresolved conflict, or changed since the version was chosen. Get the version id from `swarmfile blame`.

Argument / Flag	Description
<path>	Mount-relative (or absolute) path to the file.
--at <id>	The version's history id — the <code>#N</code> column from <code>swarmfile blame <path></code> .

```
swarmfile blame shots/seqA/plate_v3.exr # find the version id
swarmfile rollback shots/seqA/plate_v3.exr --at 4182
```

describe

Name, rename, or clear a commit's message after the fact — including an auto-saved commit that landed with no message (like `git commit --amend -m`, but for any commit and non-destructive). Author-or-admin only.

Argument / Flag	Description
<ref>	The commit to describe — the full ref vocabulary (<code>482</code> / <code>#482</code> , <code>HEAD~N</code> , a branch, a tag, a changeset id, or a 64-hex hash).
-m, --message <String>	The new message.
--clear	Remove the commit's message.

Exactly one of `-m` / `--clear` is required. Alias: `reword`.

```
swarmfile describe 482 -m "v2 delivery to client"
swarmfile describe 482 --clear
```

diff

Show what changed. With no argument (or a branch name) it previews what merging a branch into its base would change (like `git diff main...`); with two commits it shows the files that changed between them (like `git diff A B`); with one commit (any ref that isn't a branch) it shows what changed from that commit to the current branch head — `git's diff <commit>` compares against the working tree, and since saves sync as they happen here, the head is the closest equivalent. Whole-file: there's no line-level diff for a binary. Read-only — it writes nothing.

Argument	Description
<code>(none)</code> / <code>[branch]</code>	Branch preview — the given branch (or this mount's) vs its base. A branch preview is only meaningful for a non-main branch (main has no base → 400).
<code><commit></code>	That commit vs the current branch head — any non-branch ref (<code>HEAD~3</code> , a tag, a hash, <code>tag:x</code> , ...).
<code><A> </code>	Two-commit diff — files changed between commit A and B. Each accepts the full ref vocabulary: a number, <code>#N</code> , <code>HEAD</code> , <code>HEAD~N</code> , a branch, a tag, a changeset id, or a 64-hex hash (prefixes force a kind).

```
swarmfile diff           # this branch vs its base
swarmfile diff look-dev-b # a named branch vs its base
swarmfile diff HEAD~3 HEAD # what changed over the last three commits
swarmfile diff v1.0 v2.0   # between two tags
```

blame

Who last changed a file, plus its version history with authors (like `git blame` , but whole-file — one row per saved version, not per line, since content is opaque).

Argument	Description
<code><path></code>	Mount-relative (or absolute) path to the file.
<code>--branch <name></code>	Read this path's history on another branch (read-only) instead of the selected mount's — that branch's metadata is synced on demand, and the path must exist there.

```
swarmfile blame shots/seqA/plate_v3.exr
```

changesets list

The commit log, newest first, scoped to the branch this mount is on (or another named branch, read-only).

Flag	Description
<code>--limit <N></code>	Default <code>50</code>
<code>--offset <N></code>	Default <code>0</code>
<code>--branch <name></code>	Read another branch's log without switching mounts. Omitted = the selected mount's branch.

branch

Branch management.

Subcommand	Flags	Description
list	—	Active branches for this mount's project. Each row carries its head commit's hash (<code>headCommitHash</code> ; <code>null</code> when no head hash is recorded).
create	<name> , --from <REF> , --from-commit <seq> , --from-tag <tag>	Fork a new branch (defaults to forking from the selected mount's branch's live head, like <code>switch -c</code>). <code>--from</code> takes the full ref vocabulary: a branch forks that branch (an empty one included), a tag forks its commit, and a seq/hash/changeset id forks that commit with the source branch derived from it. <code>--from-commit</code> / <code>--from-tag</code> are the explicit forms; combining a branch <code>--from</code> with one of them is allowed and the hub refuses the fork if the point belongs to a different branch (<code>branch_mismatch</code>), while a second explicit point (a tag/seq <code>--from</code> plus <code>--from-commit</code> / <code>--from-tag</code>) is refused up front.
switch	[name]	Switch which branch this mount (<code>--mount</code> , else the current mount) works on. Omit <code>name</code> to switch back to the project's default branch (usually <code>main</code>).
archive	<name>	Soft-archive a branch (the project's default branch cannot be archived). Exit 2 on an actionable refusal: unknown branch, the project's default branch, a protected branch archived by a non-admin, a branch another active branch is based on (<code>branch_has_active_children</code>), or one still named by an open merge request (<code>branch_has_open_mr</code>). A merged/closed MR doesn't block, so <code>mr merge --delete-branch</code> still works.
merged	—	Branches with nothing left to land — a branch-vs-base diff with no changed files and no conflicts (the same computation <code>diff <branch></code> shows). The project's default branch is never listed, a branch whose diff can't be computed is excluded rather than assumed merged, and a branch an open mount is working on is marked <i>in use</i> .
prune	--yes	Archive every prunable branch <code>merged</code> lists, after naming them and confirming. Never archives a branch an open mount on this machine is on. <code>--yes</code> skips the prompt; <code>--json without --yes</code> returns a <code>confirmation_required</code> error instead. The result lists <code>archived</code> / <code>refused</code> / <code>failed</code> by name. Exits 2 if any branch was refused (protected, etc.), 1 on a transport error.
protect	<name> , --required-approvals <N>	Refuse a direct merge into this branch and every direct write to it (saves/edits/renames/moves/creates/deletes, staging/committing, rollback, conflict resolution, LFS reflect) — changes land only through an approved Merge Request . <code>--required-approvals</code> (1-10, default 1) sets this branch's approval floor — an unprotected branch otherwise requires none; a matching protection rule's own count can raise it further (the effective count is the MAX of both). The CI-checks-required gate (<code>--require-checks-to-pass</code>) is a <code>protection-rule</code> -only flag, not available per-branch here — protect via a single-branch (no-wildcard) rule instead if you need it without a pattern. Org-admin-gated.
unprotect	<name>	Clear a branch's protected flag. Org-admin-gated, same as <code>protect</code> .
set-commit-mode	<name> , [mode]	Lock (or unlock) this branch to a commit mode — overrides the project's own lock, if any, while set. <code>mode</code> is <code>auto-commit</code> , <code>staged</code> , or omitted/ <code>unset</code> to clear it back to unenforced-at-branch-scope (the project's lock, if any, still applies). Already-open sessions/changelists are grandfathered: this only affects the <code>next commit</code> / <code>changelist open</code> , never one already in progress. Org-admin-gated: exit 2 on <code>forbidden</code> (403) or an unknown branch (404); any other failure is exit 1.

branch `switch` is hot — no engine restart. It's refused (exit 2) while any changelist is open; submit or cancel all of them first. With the global `--mount <id>` , it retargets **that** mount only (a mount sharing another mount's branch context can't switch on its own and is refused with `shared_branch_context` — open an independent mount instead); a branch that doesn't exist can't be switched to. `branch create` never touches the running mount either, even with `--from-commit` / `--from-tag` — creating a branch doesn't change which branch this mount is checked out to. `protect` / `unprotect` exit 2 on `forbidden` (non-admin caller), `not_found` (unknown branch), or (`protect` only) an out-of-range `--required-approvals` (400) — distinguishable business states, not a crash — and exit 1 on anything else.

```
swarmfile branch create recover --from-commit 482
swarmfile branch create recover --from-tag v1.0
```

protection-rule

Glob-pattern [branch protection](#) — additive on top of `branch protect` / `unprotect` 's per-branch flag. A rule like `release/*` protects every branch whose name matches it, including ones created after the rule was added; a pattern with no wildcards is just another way to protect one exact branch.

Subcommand	Flags	Description
list	—	Protection rules for this mount's project. Org-admin-gated, like <code>create</code> / <code>delete</code> : exit 2 on <code>forbidden</code> .
create	<pattern> , --required-approvals <N> , --require-checks-to-pass	Create a rule. <code>*</code> matches any run of characters, <code>?</code> matches exactly one. <code>--required-approvals</code> (1-10, default 1) is this rule's contribution to the approval floor for every branch it matches — see <code>branch protect</code> 's row above. <code>--require-checks-to-pass</code> refuses a merge unless every CI check that has reported a run at the request's exact current head shows its latest run passing. A check whose latest run is at an earlier commit is not carried forward as still-required; if no check has reported at the head at all, the gate fails closed rather than passing vacuously. This flag exists only here, at the rule level; there's no per-branch equivalent on <code>branch protect</code> itself. Org-admin-gated: exit 2 on <code>forbidden</code> (403), <code>invalid_pattern</code> / <code>invalid_required_approvals</code> (400), or <code>pattern_taken</code> (409); exit 1 on anything else.
delete	<pattern>	Delete a rule by its exact pattern text. Exit 2 on <code>forbidden</code> or <code>not_found</code> , same convention as <code>create</code> .
add-reviewer	<pattern> , <principal-id>	Add an auto-reviewer — every merge request opened against a branch this rule matches automatically requests a review from this principal. Additive to manually-requested reviewers, never a replacement. Idempotent. Org-admin-gated.
remove-reviewer	<pattern> , <principal-id>	Remove an auto-reviewer from a rule. A no-op if they weren't one. Org-admin-gated.
list-reviewers	<pattern>	List a rule's configured auto-reviewers. Human output shortens each id (<code>user 1a2b3c4d...</code>); <code>--json</code> carries the full ids that <code>add-reviewer</code> / <code>remove-reviewer</code> take. Org-admin-gated.

```

swarmfile protection-rule create "release/*" --required-approvals 2 --require-checks-to-pass
swarmfile protection-rule list
swarmfile protection-rule add-reviewer "release/*" <principal-id>
swarmfile protection-rule list-reviewers "release/*"
swarmfile protection-rule delete "release/*"

```

tag

A project-scoped, immutable named pointer to a commit — "branch from tag" is "branch from commit, resolved via a name." A tag also pins the commit it names against history retention, so checkout/restore to it keeps working while the tag exists; deleting the tag releases the pin. There is no move/repoint subcommand: to repoint a name, delete it and create it again. (With git access on the project, a lightweight tag can also be moved or deleted over the git remote — see [Cloning a Project with Git](#).)

Subcommand	Flags	Description
list	—	Tags for this mount's project. Each row carries the tagged commit's hash (<code>commitHash</code> ; <code>null</code> when no hash is recorded for the tagged commit).
create	<name> , [REF] , --at <ref> , --at-seq <seq> , --on-branch <branch>	Tag a commit. REF takes the full vocabulary (<code>N / #N</code> , <code>HEAD~N</code> , a branch, a tag, a changeset id, a 64-hex hash; prefixes force a kind); --at-seq N is the bare-number shorthand. Omit the ref to tag the CURRENT head of the selected mount's branch, or of --on-branch (which must name an existing branch — checked with every form, --at-seq included). --from remains a hidden alias of --on-branch . Errors if the branch has no commits yet. Needs project-wide write (see Permissions); exit 2 on that refusal (403).
delete	<name>	Delete a tag. Real delete, not soft — a tag has no archived state. Same access and exit 2 on 403 as <code>create</code> .

```

swarmfile commit -m "final render pass"
swarmfile tag create v1.0
swarmfile tag create v0.9 --at-seq 482
swarmfile tag list
swarmfile tag delete v1.0

```

ref

A thin read-only grouping over branches **and** tags — for the two questions that need both listings at once. `branch` and `tag` remain the full commands for acting on either.

Subcommand	Flags	Description
list	—	Every named pointer: branches (the current one marked <code>*</code>) and tags, together — each carrying its head/tagged commit hash.
resolve <name>	—	Say what a ref resolves to under the same resolver every other command uses — the kind (<code>branch</code> , <code>tag</code> , <code>seq</code> , <code>head</code> , <code>hash</code> , <code>changeset</code>) and commit number <code>restore</code> / <code>checkout</code> / <code>log</code> would act on — with a note (<code>alsoMatches</code> in <code>--json</code>) when the name also matches another kind: a branch wins over a same-named tag, and <code>tag:<name></code> selects the tag. Exit 1 if it doesn't resolve.

```

swarmfile ref list
swarmfile ref resolve v1.0

```

release

Publish a tag as an immutable, public release — a static, cached snapshot of that tag's tree, served anonymously for browsing and gated behind a free account for full downloads. Requires the project to already be `plaintext`-tier and `public` (every public project is plaintext-tier, on any plan — see [Billing and plans](#)); publishing enqueues a combined-archive build, so a fresh release starts `pending` and moves to `ready` (or `failed`) asynchronously. `list` / `create` / `preflight` are scoped to this mount's own project; `fetch` is not — it pulls from ANY public project by org/project slug, whether or not this engine has it mounted, since a public release needs no org membership, only a real signed-in account.

A release can also carry **assets**: named binaries attached to it without living in the project tree (no manifest entry, no combined archive, no clone history). They are served anonymously at `/public/<org>/<project>/assets/<name>` on the hub, which is how build artifacts ship alongside the tagged source; `attach` uploads one, `assets` lists them, `detach` removes one, and re-attaching the same name replaces it.

Subcommand	Flags	Description
list	—	Releases published for this mount's project, each with the <code>id</code> delete takes, its <code>status</code> (<code>pending</code> / <code>ready</code> / <code>failed</code>), and the tagged commit's <code>commitHash</code> (<code>null</code> when no hash is recorded for the tagged commit).
preflight	<tag>	Report exactly what <code>create <tag></code> would do, publishing nothing: the gates (public/ <code>plaintext</code> , your authority, verified email), the file count and size that would become public, the paths <code>.swarmfile/publicignore.yml</code> keeps out, the paths <code>.gitignore</code> / <code>.swarmfileignore</code> hide from the web browser that would <i>still</i> publish, and whether the tag's policy differs from the default branch's current one (the fixed-the-policy-but-didn't-re-tag trap). Exit 0 when the tag is publishable, 2 when it isn't (the reason is printed), 1 when the check itself failed.
create	<tag>	Publish a tag. Re-running this on an already-published tag is a no-op unless the prior attempt failed non-permanently, in which case it retries. Exit 2 on a defined refusal: not <code>plaintext</code> + <code>public</code> (409), the tag's prior publish failed permanently (409 — publish a new tag instead; a terminal <code>publicignore_invalid</code> from a committed <code>.swarmfile/publicignore.yml</code> lands here, so fix the file and publish a new tag), the tag doesn't exist (404), the project isn't provisioned on the hub yet (501), or you lack the access to publish (403 — see Permissions ; this includes a folder in the tag you're denied read on, and an unverified email).
delete	<id>	Unpublish. Real delete, not soft, and needs the same access as <code>create</code> (see Permissions). Find <code>id</code> via <code>list</code> 's output first. Exit 2 on <code>404</code> (already unpublished, or a typo'd id) or <code>403</code> (you lack the access to publish, so you can't unpublish either).
attach	<tag> <file>, --name <name>, --content-type <type>	Upload a local file as an asset of the release named by <code>tag</code> (see <code>list</code>). The file's basename is the asset name unless <code>--name</code> overrides it; <code>--content-type</code> overrides the type derived from the file name. The engine computes the SHA-256 the hub displays. Re-attaching the same name replaces it. Files up to 100 MiB stream through the hub; larger files (up to 5 GiB) upload directly to the project's storage with a presigned URL. Exit 2 on a defined refusal: a bad name (400), no authority (403), a missing release (404), too many assets (409), a file over the 5 GiB cap (413), or a project bucket that cannot accept a direct upload when the file is over 100 MiB (409 <code>presign_unavailable</code>).
assets	<tag>	List the release's attached assets (name, size, content type, SHA-256).
detach	<tag> <name>	Remove one attached asset. Exit 2 on <code>404</code> (no such asset) or <code>403</code> (no authority).
fetch	--org <slug>, --project <slug>, --tag <name>, --path <path>, --out <path>, --into <org>/<project>	Fetch one file from the release's manifest, or (if <code>--path</code> is omitted) the whole release as a combined <code>.tar</code> . <code>--org</code> / <code>--project</code> are slugs, matching the public URL shape, not internal ids. With <code>--into</code> , nothing is written to disk: the release is copied straight into that project by the engine — the same client-side copy as <code>clone</code> ... <code>--into</code> (see Version control), which seals the blocks on this machine so a private destination works — <code>--out</code> is then not required, and <code>--path</code> narrows the copy exactly as it narrows an archive. The copy runs as a long operation (<code>--no-wait</code> , <code>`swarmfile op status</code>

Browsing a published release anonymously (no account, no CLI) is a web-hosted page, not a CLI surface — the same split [share](#) describes for share links.

```
swarmfile tag create v1.0
swarmfile release preflight v1.0
swarmfile release create v1.0
swarmfile release attach v1.0 dist/widgets-1.0-macos-arm64.tar.gz
swarmfile release assets v1.0
swarmfile release detach v1.0 widgets-1.0-macos-arm64.tar.gz
swarmfile release delete re_9f2a
swarmfile release list
swarmfile release fetch --org acme --project widgets --tag v1.0 --out widgets.tar
swarmfile release fetch --org acme --project widgets --tag v1.0 --path docs/readme.md --out readme.md
swarmfile release fetch --org acme --project widgets --tag v1.0 --into myorg/private-widgets
```

merge

Merge this mount's branch into the branch it was forked from — its parent (usually `main`, but a branch forked from a non-main branch merges back into that one). The target is implicit: a branch already knows its own parent, so there's no target flag. A merge that would change an entry with an unresolved conflict is refused with `conflict_unresolved` — resolve it first, then re-run.

Flag	Description
--from <branch>	Merge FROM this branch instead of the selected mount's own — land a branch you aren't standing on. The target is still the source's parent (main, or whatever it was forked from).
--record-conflicts	Record <code>entry_conflicts</code> rows for interactive resolution instead of only listing them on conflict.
--diff3	On a merge conflict, batch-resolve every text-file content conflict with the engine's three-way merge before giving up, then retry the merge once. Never lands a partial merge on its own. Overrides a project's <code>.swarmfile/merge.yml</code> <code>diff3: false</code> . Mutually exclusive with <code>--no-diff3</code> .
--no-diff3	Force diff3 auto-resolve off for this invocation, even if the project's <code>.swarmfile/merge.yml</code> declares <code>diff3: true</code> . Mutually exclusive with <code>--diff3</code> .
--dry-run	Print what the merge would change — added, modified, deleted, renamed, and conflicting files, the same read-only diff <code>swarmfile diff <branch></code> shows — and stop. Nothing is staged and no merge lock is taken.
--wait	Large merges apply in the background; the command otherwise returns as soon as the merge is accepted. Block and poll until the apply finishes, printing each phase. Gives up after ten minutes with exit 2 and explains how to resume or roll the merge back.

```
swarmfile branch create feature-x
swarmfile branch switch feature-x
# ...work, commit...
swarmfile branch switch
swarmfile merge --record-conflicts
```

mr

Open, list, inspect, review, and land [Merge Requests](#) — the propose → review → approve → merge gate that sits on top of `merge` above.

Subcommand	Flags	Description
<code>create</code>	<code>--title <title></code> , <code>--from <branch></code> , <code>-b <branch></code> , <code>--body <text></code> , <code>--draft</code>	Open a merge request from <code>--from</code> into its base branch. <code>--from</code> defaults to whatever branch this mount is currently on. Prints the branch-vs-base diff preview (same as <code>diff <branch></code>) before opening it, so you see what a reviewer will see. <code>--draft</code> opens it as a <code>draft</code> — reviewable and commentable, but <code>mr merge</code> refuses it, and the usual "new merge request" notification never fires for it at all (not deferred — marking it ready later doesn't retroactively send the one it skipped).
<code>list</code>	<code>--status <open merged closed all></code> (default <code>open</code>), <code>--limit <N></code> , <code>--label <label-id></code> , <code>--assigned-to <principal-id></code>	Merge requests for this mount's project. <code>--limit</code> is a hint the hub caps at 500 regardless. <code>--label</code> / <code>--assigned-to</code> filter server-side — see <code>label list</code> / <code>assignees</code> , <code>requested reviewers</code> , and <code>labels</code> for where the ids come from. <code>--assigned-to</code> matches the <code>assignee</code> role specifically, not a requested reviewer.
<code>status</code>	<code><number></code>	Branch names, who opened it, approval state, per-reviewer verdicts, assignees/requested-reviewers/labels, and changed-file count for one merge request, by its display number (the "3" in "MR-003"). Exit code doubles as a CI gate: <code>0</code> approved, <code>2</code> open but not yet approved, <code>1</code> any other error.
<code>diff</code>	<code><number></code>	The merge request's live diff — the files it would change (A/M/D/R) plus any conflicts, rendered exactly like <code>diff</code> above. Read-only, recomputed fresh by the hub on each call. If the MR's source branch was archived/deleted the hub has no diff and this says so, rather than printing an empty list that reads as "no changes."
<code>approve</code>	<code><number></code> , <code>--note <text></code>	Approve a merge request. Refused (<code>self_review</code>) if you opened it yourself.
<code>request-changes</code>	<code><number></code> , <code>--note <text></code>	Request changes on a merge request — blocks <code>merge</code> below until a later review approves it instead. Same self-review restriction as <code>approve</code> .
<code>review-comment</code>	<code><number></code> , <code>--note <text></code>	Leave a non-blocking "Comment" review — shows up in the review history but never counts toward the approval floor and never blocks <code>merge</code> . Unlike <code>approve</code> / <code>request-changes</code> , allowed on your own merge request. Distinct from <code>comment</code> / <code>comments</code> below, which post to the discussion thread, not the formal review history.
<code>ready</code>	<code><number></code>	Mark a <code>draft</code> merge request ready for review — lifts <code>merge</code> 's draft refusal, and notifies the merge request's current assignees and requested reviewers (<code>mr_ready</code>). Opening it as a draft suppressed the one-time "new merge request" notification, which is not re-sent. One-directional: there's no way back to draft. Refused (<code>not_draft</code>) if it isn't currently a draft, or (<code>invalid_state</code>) if it isn't open.
<code>merge</code>	<code><number></code> , <code>--record-conflicts</code> , <code>--delete-branch</code> , <code>--diff3</code> , <code>--no-diff3</code>	Land an approved merge request — same conflict detection as <code>merge</code> above, including the <code>--diff3</code> / <code>--no-diff3</code> override flags. Refused with <code>not_approved</code> unless enough reviewers have approved (zero on an unprotected branch; a <code>protected target branch or rule</code> raises the floor, 1-10), nobody's latest verdict is "request changes," and (if the target requires it) CI checks pass. Refused with <code>is_draft</code> while the merge request is still a draft — run <code>mr ready</code> first. <code>--delete-branch</code> archives the source branch once the merge has actually landed — opt-in, off by default; a failed archive (e.g. the branch got protected in the meantime) never fails the merge itself, just prints a warning.
<code>close</code>	<code><number></code>	Close a merge request without merging it.
<code>reopen</code>	<code><number></code>	Reopen a closed merge request. Refused (<code>invalid_state</code>) unless it's currently closed — a merged merge request can never be reopened.
<code>comment</code>	<code><number></code> , <code>--message <text></code>	Post a comment to a merge request's discussion thread — the same thread the web dashboard's merge request page shows.
<code>comments</code>	<code><number></code>	List comments on a merge request's discussion thread.
<code>assign</code>	<code><number></code> , <code><principal-id></code>	Assign a principal (user or group id — see <code>admin nodes</code> /the web dashboard's people picker for ids) to land this merge request. Idempotent: assigning someone already assigned is a no-op, not an error.
<code>unassign</code>	<code><number></code> , <code><principal-id></code>	Unassign a principal. A no-op if they weren't assigned.
<code>request-review</code>	<code><number></code> , <code><principal-id></code>	Request a review from a principal — distinct from <code>assign</code> (who lands it). Idempotent.
<code>unrequest-review</code>	<code><number></code> , <code><principal-id></code>	Withdraw a review request. A no-op if none was pending.
<code>label</code>	<code><number></code> , <code><label-id></code>	Attach a label (from <code>label list</code> below) to this merge request. Idempotent.
<code>unlabel</code>	<code><number></code> , <code><label-id></code>	Detach a label. A no-op if it wasn't attached.

```

swarmfile mr create --title "Warmer facade material"
swarmfile mr create --title "WIP: new roof profile" --draft
swarmfile mr ready 4
swarmfile mr list
swarmfile mr list --label <label-id> --assigned-to <principal-id>
swarmfile mr status 3
swarmfile mr diff 3
swarmfile mr approve 3
swarmfile mr request-changes 3 --note "needs another pass"
swarmfile mr merge 3
swarmfile mr close 3
swarmfile mr reopen 3
swarmfile mr comment 3 --message "Looks good once the fascia texture is swapped."
swarmfile mr assign 3 <principal-id>
swarmfile mr request-review 3 <principal-id>
swarmfile mr label 3 <label-id>

```

approve / request-changes / merge / close put every review action a terminal command away. The hub refuses a review from the merge request's own creator; an approval from a second credential counts, so treat review-capable credentials like any other approval authority in your org. `branch protect --required-approvals` raises the bar for a unilateral change.

Label

A project's label registry — the tags `mr label / mr list --label` above attach to and filter by. See [assignees, requested reviewers, and labels](#).

Subcommand	Flags	Description
list	—	This mount's project's label registry.
create	<name>, --color <#RRGGBB>	Define a new label. Org admin/owner only.
update	<label-id>, --name <name>, --color <#RRGGBB>	Rename and/or recolor a label in place — the alternative to delete-and-recreate, which would drop every existing merge request's attachment. Both flags optional; an omitted one is left unchanged. Org admin/owner only.
delete	<label-id>	Delete a label — cascades off every merge request it was attached to. Org admin/owner only.

```

swarmfile label list
swarmfile label create bug --color "#FF5733"
swarmfile label update <label-id> --name defect --color "#000000"
swarmfile label delete <label-id>

```

rfi

RFIs and Submittals — the AEC request-for-information / approval workflow, the same feature the dashboard's RFI panel drives. Every command here (other than `create / list / search / by-file`) takes an RFI, and you can name it either way: the friendly display number (`RFI-003` , case-insensitive — resolved through the mounted project) or its internal `id` from `rfi create` 's or `rfi list` 's output. Resolving a number needs a mounted project and the same read access the detail view has.

Subcommand	Flags	Description
<code>create</code> <code><path></code>	<code>--title</code> , <code>--body</code> (required); <code>--priority</code> <code><low normal high urgent></code> ; <code>--due-at</code> <code><unix-s></code> ; <code>--assigned-to</code> <code><id></code> ; <code>--reviewer</code> <code><id></code> (repeatable); <code>--kind</code> <code><rfi submittal></code> (default <code>rfi</code>); <code>--draft</code>	Open an RFI (or, <code>--kind submittal</code> , a Submittal) rooted at a mount-relative folder — same path resolution as <code>comment</code> / <code>watch</code> . Needs <code>write</code> there. <code>--draft</code> creates it unsubmitted; call <code>submit</code> when ready.
<code>list</code>	<code>--status</code> , <code>--priority</code> , <code>--ball-in-court</code> <code><originator reviewer both></code> , <code>--kind</code> , <code>--assigned-to</code> , <code>--due-before</code> <code><unix-s></code> , <code>--outstanding</code> , <code>--limit</code> , <code>--offset</code>	RFIs/Submittals matching the filters. <code>--outstanding</code> narrows to the same still-needs-attention set (excludes <code>closed</code> / <code>void</code> / <code>draft</code>) the file-browser's badge uses.
<code>status</code> <code><id></code>	—	Full detail, including decision history. Exit <code>0</code> if <code>answered</code> / <code>closed</code> (someone has responded), exit <code>2</code> for every other state (<code>draft</code> / <code>open</code> / <code>in_review</code> / <code>void</code>) — a CI/pre-merge gate can check "has this been answered."
<code>update</code> <code><id></code>	<code>--title</code> , <code>--body</code> , <code>--priority</code> , <code>--due-at</code> , <code>--assigned-to</code> , <code>--ball-in-court</code>	Change mutable fields — every flag optional, only the ones given change. Originator-or-owner only.
<code>submit</code> <code><id></code>	—	<code>draft</code> → <code>open</code> . Originator-or-owner only.
<code>acknowledge</code> <code><id></code>	—	<code>open</code> → <code>in_review</code> — a reviewer marking they've seen it. Assigned-reviewers only.
<code>close</code> <code><id></code>	—	<code>answered</code> → <code>closed</code> . Originator-or-owner only.
<code>void</code> <code><id></code>	—	<code>any</code> → <code>void</code> — abandon it (terminal, leaves the number gap). Originator-or-owner only.
<code>decide</code> <code><id></code> <code><decision></code>	<code>-m</code> , <code>--message</code> <code><text></code> (required); <code>--signature-kind</code> <code><click drawn></code> ; <code>--signature-blob</code> ; <code>--comment-id</code>	Record a decision — an RFI's answer, or a Submittal's review outcome. <code>decision</code> is one of <code>approved</code> , <code>rejected</code> , <code>revisions_requested</code> , <code>answered</code> , <code>commented</code> , <code>approved_rev_a</code> , <code>approved_rev_b</code> , <code>approved_as_noted</code> , <code>rejected_as_noted</code> — the last four are Submittal-only; the hub doesn't itself reject a mismatched kind, so check <code>status</code> <code><id></code> 's <code>kind</code> first. Comment-or-above ACL.
<code>decisions</code> <code><id></code>	<code>--limit</code> , <code>--offset</code>	The decision audit trail, oldest first.
<code>assign</code> <code><id></code>	<code>--add</code> <code><principal-id[:role]></code> (repeatable), <code>--remove</code> <code><principal-id></code> (repeatable)	Add/update or remove reviewer-role assignments. <code>role</code> is <code>reviewer</code> (default) / <code>observer</code> / <code>originator</code> . Admin-gated.
<code>attachments</code> <code><id></code>	—	Files attached under the RFI's root.
<code>search</code> <code><query></code>	<code>--project-id</code> , <code>--limit</code>	FTS5 search over title/body, ACL-filtered to what you can read. Defaults to the mounted project; pass <code>--project-id</code> to search a project this engine isn't mounted on.
<code>by-file</code> <code><path></code>	—	Resolve the RFI (if any) that owns a file — the CLI equivalent of the dashboard's right-click "Open RFI thread."

```

swarmfile rfi create ./Projects/tower-a/foundations --title "Slab thickness at grid C4" --body "Drawing S-201 shows 300mm, spec calls for 350mm
- which governs?" --priority high --reviewer user_88
swarmfile rfi list --outstanding
swarmfile rfi status rfi_9f2a
swarmfile rfi decide rfi_9f2a answered -m "350mm per spec - drawing will be revised"
swarmfile rfi status rfi_9f2a # now exits 0

```

ec-placement

Deliberate erasure-coded shard placement across LAN peers.

Subcommand	Description
<code>status</code>	Whether this mount currently has deliberate LAN shard placement on.
<code>enable</code>	Turn on deliberate LAN shard placement.
<code>disable</code>	Turn off deliberate LAN shard placement.

`ec-placement enable` / `disable` restart the engine to take effect.

Placement spreads shards across your **LAN peers**, which are normally found by mDNS. If your network blocks mDNS multicast (common on corporate/VLAN'd networks), the office won't be discovered and every shard falls back to this machine — set `lan_from_office` (see [Engine Config File](#)) so same-office peers count as LAN.

shards <path>

Where a file's erasure-coded shards live across the office — for each piece, the machines it's assigned to and whether this one holds it. Requires `ec-placement enable` and a live P2P fabric. The command opens with a per-file summary: how many chunks and shards the file has, whether this machine holds enough shards (`k` per chunk) to reconstruct it alone, and whether placement actually assigned every shard to `K` peers — so "is this file protected in the office, and from where" is one line instead of counting rows.

Day-to-day

status

This mount's engine, peers, and settings at a glance. It also prints the mounted branch's head commit pin when one is known (`head commit: <12-char> (<full hash>)`) so a script can capture the exact commit this workspace is on.

`pending uploads` counts this mount's project. When another project still has saves uploading (one you switched away from keeps uploading in the background), the line adds (N more in other projects) and one indented line per such project follows, named when the engine has seen the project's name and by id otherwise, with any saves held there by a refusal. `--json` carries the same as `pendingUploadsAllProjects` and `pendingUploadsByProject`.

`--changes` swaps that for a Git-status-like summary of what's outstanding on this mount: the current changelist's staged files and any blocked conflicts. There's no "modified but unsaved" state to show — saves either stage into the changelist or commit immediately — so "staged" plus "conflicted" is the whole picture.

```
swarmfile status          # engine, peers, settings
swarmfile status --changes # staged files + conflicts
```

stats

Where this machine's read bytes actually came from — per file and in total: **this computer**, **a colleague's computer**, or **the cloud** — plus first-read timings. The engine counts each block **once per file per counter window**, at the moment that file is first served it (a cache hit is "this computer"; a LAN/WAN peer transfer is "colleague"; a fetch from cloud storage is "cloud"), including read-ahead, so the cloud and colleague figures are the bytes really downloaded and re-reading or scrubbing a timeline never inflates them. Counters are per project and per engine process; `--reset` starts a fresh window (use it between takes so each recording's numbers stand alone — it zeroes both this project's breakdown and the engine's process-wide bandwidth counters), and the tray's **Data sources** panel shows the same counters live with its own reset button.

`--json` prints the full evidence report — engine version, machine, OS/arch, the mount path/branch, file sizes, per-file source split, read counts and first-read latency — one object you can paste into an evidence log or a journalist can reproduce against. `--csv` prints the same as spreadsheet rows (a header, a summary row and one row per file, each carrying date/version/machine/project) for the evidence log.

```
swarmfile stats          # human table
swarmfile stats --json   # the evidence report
swarmfile stats --csv    # spreadsheet rows, one take per machine
swarmfile stats --reset  # zero the counters after printing
```

demo reset

`swarmfile demo reset [--project <id>] [--force] [--yes]` — start a demo take cold on **this machine**: empty the project's local cache (the same whole-project "Free up space" sweep, including "available offline" marks) and zero the `stats` counters — this project's window and the process-wide bandwidth counters, so every surface the recording shows starts at zero. Run it on every machine a recording uses; a second take that quietly read from the cache is the failure this exists to prevent.

Files with saves that have not committed yet are **refused** (409) — their bytes may be the only copy — unless `-f / --force`, which proceeds while still keeping them. Cached content re-downloads on demand and nothing is deleted at the hub. The confirmation prompt is skipped with `-y / --yes`.

```
swarmfile demo reset --yes
swarmfile demo reset --project prj_... --force
```

demo offline

`swarmfile demo offline on [--include-lan] [--for <duration>]`, `swarmfile demo offline off`, `swarmfile demo offline status [--json]` — cut **this machine's** network for real, on cue, for a recording of "the internet drops mid-work". It installs host firewall rules (macOS pf, Linux nftables or iptables, Windows Firewall), so the engine sees exactly what a dead office uplink looks like; it is not a simulation, and it works with no engine running.

- **The office LAN stays up by default.** Everything is blocked except this machine itself, private addresses (10/8, 172.16/12, 192.168/16), link-local, IPv6 unique-local and multicast — so office computers keep serving each other while the cloud is out of reach. `--include-lan` blocks the LAN too (everything but this machine); a remote shell to it drops.
- **Existing connections are cut too**, including the engine's long-lived connection to the hub — not just new ones.
- **It always switches itself back off.** `on` arms an automatic `off` after 30 minutes (`--for 90s`, `--for 10m`, `--for 1h30m`; at most 24h) that survives the command, the terminal, and on Windows a restart. `off` ends it sooner and cancels the timer; running it when nothing is on is harmless.
- **Needs administrator rights** for `on` and `off`: on macOS and Linux run it with `sudo`; on Windows, from an elevated terminal. Without them it exits `2` and prints the exact command to run. `status` needs none; `--json` reports `offline`, `scope`, `lanStaysUp`, `since`, `restoreAt` and `remainingSecs` for scripts.

```
sudo swarmfile demo offline on --for 10m # internet off, office LAN up
swarmfile demo offline status
sudo swarmfile demo offline off          # back online
```

op

Some commands do work proportional to your data — `merge`, `commit`, `checkout`, `revert`, `materialize`, `branch create/archive`, `mr merge`, `trash restore/purge`, `project delete`, `ignore-clean`, `hydrate free-space`, and opening or closing a mount. They run as **operations** in the engine: the command waits and prints the result as usual, however long that takes, but the work no longer depends on the terminal staying open.

- **Ctrl-C detaches** instead of stopping anything: the command exits `3` and prints the operation's id, and the work carries on.
- `--no-wait` returns the id at once (exit `3`).
- An operation's outcome is kept, so you can collect it later — even after the engine restarts. One still running when the engine stopped is reported as interrupted.
- When an operation follows a hub background job — an over-cap folder delete/restore/purge or `project delete` (see [Trash, History & Rollback](#)) — `op status` also shows its `done / total` progress while it runs.

Subcommand	Description
<code>op list</code>	Recent operations, newest first.
<code>op status <id></code>	One operation's state and, once finished, its result. Exits <code>0</code> succeeded, <code>1</code> failed, <code>2</code> canceled, <code>3</code> still running.
<code>op wait <id></code>	Wait for it to finish and print its result; exits the way the original command would have.
<code>op cancel <id></code>	Stop an operation that can safely stop part-way (<code>materialize</code> , <code>hydrate free-space</code> , <code>git backfill-oids</code> , <code>git index</code> , and <code>checkout / revert</code> while still staging). A cancel that arrives once the staged apply has begun is declined (<code>too_late</code>) and the operation ends with the commit's real outcome. Commits, merges and deletes complete or fail on their own and refuse (<code>not_cancellable</code>).

```
swarmfile --no-wait merge --from feature-x # prints the operation id, exit 3
swarmfile op wait 1b2c...                 # the merge's own result, when it lands
```

mounts

Manage this engine process's concurrent mount points. One engine can hold several full mounts open at once — of the same project, or (same org only) a genuinely different one via `--project` — each with its own changelists, metadata, and lock manager — for running multiple isolated mounts without paying git's one-worktree-per-branch tax (the motivating case: several AI coding agents, each working against the same project without stepping on each other's saves). Most installs only ever run the one boot mount; this is for when you deliberately want more than one. Multiple live mounts are supported on every platform, with a bounded number per machine (18 on macOS): an `open` fails with `mount_failed` once that limit is reached — close a mount to free a slot. For the headless, per-agent recipe (API key, labeled mounts, shared cache), see [Coding Agents](#).

Subcommand	Flags	Description
<code>list</code>	—	Every mount this engine currently has open, (<code>current</code>) marking the one id-less commands (and any command run without <code>--mount</code>) target.
<code>open <mount_point></code>	<code>--label <text></code> , <code>--set-current</code> , <code>--branch <name></code> , <code>--project <id></code> , <code>--exclusive</code>	Open a new mount at a filesystem path, a drive letter (Windows: <code>P:</code>), or <code>auto</code> — the next free drive letter on Windows, or a new folder beside the boot mount on macOS and Linux, named after the label or branch. On Windows a folder must be empty or not exist yet (Windows creates the mount folder itself); a folder with anything in it is refused with <code>mount_point_not_empty</code> , and a drive letter something else holds with <code>drive_letter_in_use</code> , naming what holds it. The command returns once the drive has attached, or with the reason it couldn't (<code>mount_failed</code>) — a mount that fails to attach is not left registered. It is a full, symmetric mount, equally capable to the boot mount, not a lightweight second view of it. Does not become current unless <code>--set-current</code> is given, so opening one never retargets a command already running elsewhere against the existing mount. <code>--branch</code> is optional; omit it to mount the project's currently-active branch — but a mount opened without <code>--branch</code> shares the engine's boot branch context and can't be switched independently (<code>mounts branch</code> refuses it with <code>shared_branch_context</code>). Naming a <i>different</i> branch opens a divergent (<code>project</code> , <code>branch</code>) scope: the mount works that branch while every other mount stays on its own, and it can be switched later with <code>mounts branch</code> . <code>--project <id></code> opens a mount on a DIFFERENT project than this engine's active one (same org only); it combines with <code>--branch</code> , and a cross-project mount pinned to a branch gets that branch's own live-update feed. <code>--exclusive</code> opts the mount into exclusive-write mode : for as long as a file is open for writing through it, a <i>sibling</i> mount opened with the same flag on the same project+branch has its write-open refused (the app sees <code>EAGAIN</code> , the same refusal another machine's lock produces) — a rename, delete, or atomic-save replace of that file is refused too — instead of both editing the file and letting whichever closes last win. A mount that does not set the flag neither takes nor honors claims, so existing multi-mount setups are unchanged; open every agent mount with it when several agents share one tree. The flag is fixed for the mount's lifetime and shown by <code>mounts list / mounts status</code> as <code>exclusive</code> .
<code>close <id></code>	<code>--force</code>	Close a mount. Refused (exit 2, <code>changelists_open</code>) if it has open changelists, unless <code>--force</code> is given — never auto-cancels staged work. Refused (exit 2, <code>is_default_mount</code>) for the boot/default mount specifically: closing it shuts down the whole engine, not just one mount, so it isn't reachable through this command at all — quit the engine itself (the Desktop App's Quit, or stopping the <code>swarmfile-engine</code> process) if that's what you actually want.
<code>current <id></code>	—	Switch which mount id-less commands target, from then on.
<code>branch <id> [branch]</code>	<code>--wait</code>	Switch one mount's branch independently of every other mount (omitted = the project's default branch). Only works on a mount that was opened with its own <code>--branch</code> ; a mount sharing the engine's boot branch context is refused with <code>shared_branch_context</code> . <code>--wait</code> blocks until the (asynchronous) switch lands or fails.
<code>status <id></code>	—	One mount's own detail — the same fields <code>list</code> shows, for just this id, without scanning the whole list.
<code>scratch-dir <id></code>	—	That mount's project's machine-local scratch/dependency-cache directory, plus whatever tool-cache redirects its <code>.swarmfile/cache.yml</code> declares (see <code>scratch</code> below). Creates the directory — and every declared subdirectory — if it doesn't already exist.

The global `--mount <id>` flag (see "Global flags" above) is how every OTHER subcommand targets a specific mount instead of `current` — e.g. `swarmfile --mount mount-abc123 status` runs `status` against that mount regardless of which one is current. It has no effect on `mounts` subcommands themselves, which already take the id as their own argument, nor on `workspace switch` (below), which retargets the whole engine's active workspace; combining the flag with either prints a warning and is otherwise a no-op for the flag.

```

swarmfile mounts open ~/mount-2 --label "Agent 2" --set-current
swarmfile mounts open auto --label agent-3 --branch feature-x # Windows: next free drive letter
swarmfile mounts open ~/agent-3 --exclusive # exclusive-write mode for agent fleets
swarmfile mounts open ~/src-only --sparse "src/**" # a sparse mount - only src/ is visible
swarmfile mounts open ~/no-vendor --exclude "vendor/**" # everything except vendor/
swarmfile mounts open ~/other-project --project proj_abc123 # different project, same org
swarmfile mounts list
swarmfile mounts status mount-abc123
swarmfile mounts current mount-abc123
swarmfile --mount mount-abc123 status # target it without switching current
swarmfile mounts close mount-abc123

```

Sparse mount. `mounts open --sparse <glob>` (repeatable) exposes only the paths matching those include globs — the same dialect as `materialize --sparse` (a glob matching a directory includes its subtree, a leading `!` excludes an earlier match, a bare name matches at any depth). `--exclude <glob>` is sugar for the `!` negations; with no `--sparse` it means "everything except". An excluded path is invisible at that drive: no listing entry, an open fails as if it weren't there, and a write-create/mkdir/rename onto it is refused. Empty = the whole branch, and the scope is per mount (a sibling mount is unaffected); a scoped mount's offline (`hydrate /Pack & Go`) set matches the scope.

workspace

Show or change the org/project this engine is scoped to — the CLI equivalent of the Desktop App's org and project pickers. A same-org project switch usually takes the hot path (no restart), and an org switch does too while the mount is live; switching to a project the running mount can't be re-scoped to, or switching with no live mount to work with, drains and restarts the engine, so the command may briefly drop its connection. A target project whose key hasn't been released yet doesn't force that restart — the switch waits for the grant and completes in place. A same-org switch *with creates still queued* can also take that restart fallback — queued creates survive it and land in the project they were created under. Refused (exit 2) while a changelist is open, or while another switch is already running.

Subcommand	Flags	Description
<code>status</code>	—	The current org and active project for this mount.
<code>list</code>	—	Every org this account can reach, plus every project in the current org, with <code>(active)</code> marking the current pair. Use these ids with <code>switch</code> .
<code>switch</code>	<code>--org <id></code> , <code>--project <id></code> , <code>-y / --yes</code> (<code>--force</code>), <code>--park</code>	Move the engine to that org, and optionally that project in it. <code>--org</code> is required; omitting <code>--project</code> clears the stored project so the engine adopts the target org's default. A switch is refused (exit 2) while a changelist is open — <code>--park</code> sets the mount's open changelists aside first so the switch can proceed; <code>--yes</code> skips the confirmation prompt (which names staged files and queued/failed writes).

```

swarmfile workspace status
swarmfile workspace list
swarmfile workspace switch --org org_abc --project proj_9f2a
swarmfile workspace switch --org org_abc # adopt that org's default project

```

`workspace switch` changes the whole engine's active workspace, not one mount, so the global `--mount` flag has no effect on it (the CLI prints a note if you combine them).

conflicts

Lists files that are blocked because the same file was changed in two places and the two versions haven't been reconciled yet, and resolves them.

While a file is in that state, writes to it are refused, and an application sees a plain I/O error — a filesystem has no way to explain a cross-machine change. This command is where the reason lives.

Subcommand	Flags	Description
<code>list</code> (default — bare <code>conflicts</code> is the same)	—	Files blocked by an unresolved conflict, by path. A row on another branch (a branch-merge conflict is recorded under the merge TARGET's entry id) prints the branch name, or <code>(not on this mount's branch)</code> when the hub couldn't name it.
<code>resolve <path></code>	<code>--winner</code> <code>local remote keep-both</code>	Accept one side non-interactively, without running any tool.
<code>resolve <path></code>	<code>--tool [name]</code>	Resolve via an external mergetool speaking git's own <code>mergetool</code> protocol. Mutually exclusive with <code>--winner</code> and <code>--auto</code> .
<code>resolve <path></code>	<code>--auto</code>	Run the engine's three-way diff3 auto-merge and accept the result if it's clean — no external tool, no <code>gitconfig</code> needed. Exits 2 when a real conflict remains (a text collision, or a binary).
<code>resolve --all [--auto --winner S]</code>	—	Attempt every conflicted file on this mount's branch, printing one summary at the end. Requires <code>--auto</code> (engine merge) or <code>--winner <local remote keep-both></code> (accept that side everywhere). Rows on other branches are skipped and counted unless <code>--include-other-branches</code> is given — pass it when landing a merge (the conflict is recorded under the target branch), or use <code>merge --diff3</code> . Exits 2 if any file still needs a human (auto only), 1 if any request failed.
<code>resolve --all --include-other-branches</code>	—	With <code>--all</code> , also act on conflicts the hub reports for other branches.

```
$ swarmfile conflicts
1 file(s) blocked by an unresolved conflict:

/Projects/plan.dwg
```

Writes to these files are refused until the conflict is resolved.
 Resolve with `'swarmfile conflicts resolve <path> --auto', '--winner local|remote|keep-both', or '--tool', in the Desktop App, or in the web dashboard.`

`--winner` picks a side directly — same three outcomes as the Desktop App/web picker (keep mine, keep theirs, keep both), just from the CLI:

```
swarmfile conflicts resolve /Projects/plan.dwg --winner local
```

For a conflict recorded by `merge --record-conflicts`, `local` is the incoming branch's version and `remote` is the version already on the branch being merged into — see [Branches & Merging](#).

When `--winner remote` keeps a version different from yours, your edit is preserved in the file's history (it shows in `blame` and the History views) and the command prints a note saying so; `swarmfile rollback --at <id>` can bring it back.

`--auto` runs the engine's own three-way diff3 merge and submits it when it comes back clean. It needs nothing installed and no `gitconfig`, which makes it the resolution path to reach for from a pre-flight or CI script:

```
swarmfile conflicts resolve /Projects/plan.dwg --auto
```

It exits **2** (not 1) when auto-merge can't produce a clean result — a genuine three-way text collision, or a binary / non-diff3-eligible file — so a script can tell "a human must choose" apart from "the engine was unreachable" and fall back to `--winner` or `--tool`.

Add `--all` to sweep every conflicted file in the project in one pass and clear the text ones automatically, leaving only the files that genuinely need a person:

```
swarmfile conflicts resolve --all --auto
```

`--all` also accepts `--winner`, to accept one side across every conflicted file in one pass (useful when a whole branch should keep yours, or theirs):

```
swarmfile conflicts resolve --all --winner local
```

`--tool [name]` resolves a whole-file conflict — including binary files the Desktop App/web picker's "try automatic merge" can't touch — via an external mergetool speaking git's own protocol: `$BASE / $LOCAL / $REMOTE / $MERGED` temp-file substitution, same as `git mergetool`. This is the one resolution path that's CLI-only; it needs a real terminal/GUI session to run the tool in, which only the foreground CLI process has.

Resolution order is: swarmfile's own `mergetool_cmd` config (see [the config reference](#)) if you've set one — it wins over `gitconfig` without even reading it — else your `git config's [merge] tool + [mergetool "<name>"] cmd`, so anyone with `git mergetool` already configured and no swarmfile-specific override gets it for free. Name one explicitly to bypass both:

```
swarmfile conflicts resolve /Projects/plan.dwg --tool          # configured default
swarmfile conflicts resolve /Projects/plan.dwg --tool kdiff3  # a specific tool
```

A nonzero exit from the tool aborts before anything uploads — nothing is resolved, and the CLI tells you where the scratch files (`$BASE / $LOCAL / $REMOTE / $MERGED`) are for a manual look or retry.

throttle

Bandwidth throttling — applies immediately, no restart.

Subcommand	Flags	Description
status	—	Whether throttling is on, the rate in effect right now, and why (office or off hours, and that window's percentage of the site link).
enable	<code>--download-mbps <N></code> , <code>--upload-mbps <N></code> , <code>--office-hours-start <HH:MM></code> , <code>--office-hours-end <HH:MM></code> , <code>--office-hours-wan-pct <0-100></code> , <code>--off-hours-wan-pct <0-100></code>	Turn throttling on. Omit a flag to leave that value at whatever it already resolves to.
disable	—	Turn throttling off.

The `--office-hours-* / --off-hours-*` flags shape a time-of-day schedule: WAN traffic is capped to a percentage of the site link during the office-hours window (`--office-hours-start / --office-hours-end`, HH:MM local time) and to a different percentage outside it, so background sync can back off while people are working and open up overnight. The window's start and end must be different times: to use one rate all day, set both percentages to the same value. These same values are file-settable as `office_hours_start / office_hours_end / office_hours_wan_pct / off_hours_wan_pct` (see [Engine Config File](#)).

```
swarmfile throttle enable --office-hours-start 08:00 --office-hours-end 18:00 \
--office-hours-wan-pct 20 --off-hours-wan-pct 80
```

What the cap covers. Bulk background transfer: uploads, `hydrate`, and `fetch`. A read an application is actually waiting on — opening a file through the drive — is **not** throttled, so a cap set to protect the office link can't turn somebody's file open into a stall. The practical consequence is that a large `hydrate` or `fetch` is bounded by this setting: both are pre-fetches filling the cache ahead of use, not somebody sitting in front of a spinner.

cache

Local disk cache size cap — applies immediately, no restart. This is the block cache (the project's own synced content); for machine-local package-manager/build caches shared across mounts, see `run` and `scratch` below — a different, unrelated cache.

Subcommand	Flags	Description
status	—	Current cache cap and usage.
set	<gib: u64>	Set the local disk cache cap, in GiB.
reclaim	--dry-run , --yes	Remove the orphaned pre-project-scoping flat block cache (its blocks re-download on demand). <code>--dry-run</code> reports without deleting; otherwise it asks first (<code>--yes</code> skips).
resync	—	Re-fetch this project's hub change feed from the beginning — a repair for a mount whose feed cursor advanced past a change it never applied (an entry missing on that machine alone while the hub serves it correctly). Applies are idempotent; local rows and cached bytes are kept, and the walk runs in the background with the engine retrying until it catches up.

run

Run a command with this project's declared tool caches (see `.swarmfile/cache.yml`, under `scratch` below) redirected into environment variables — e.g. `swarmfile run -- pnpm install` points pnpm's store at a directory every mount of this project, on this machine, shares, so a second mount (or a second AI agent's mount) doesn't pay to re-download what the first one already fetched. Put `—` before the command so its own flags aren't parsed as swarmfile's.

Always safe to prefix onto anything: with no `.swarmfile/cache.yml`, a project that declares none, a config with a bad entry, or even an unreachable engine, `run` just warns on stderr and runs the command completely unmodified — it never refuses to run the command over a cache-resolution problem. The wrapped command's stdio and exit code are passed straight through. For wiring several agent mounts onto one shared cache, see [Coding Agents](#).

```
swarmfile run -- pnpm install
swarmfile run -- cargo build --release
```

scratch

This machine's per-project scratch/dependency caches — the machine-local, non-versioned directories `run` (above) redirects tool caches into. One directory per project, shared by every mount of that project on this machine regardless of branch, the same way a developer's real `~/.cargo/registry` isn't per-checkout either. The worked example (agent mounts sharing one cache, with `gc`) is in [Coding Agents](#).

Subcommand	Flags	Description
list	—	Every project's scratch cache on this machine: size on disk, how long since it was last used, and whether that project is currently open in this engine.
clear <name>	--force	Delete ONE project's scratch cache outright, by the name <code>list</code> shows. Refuses if that project is currently open, unless <code>--force</code> .
gc	--older-than-days <N> (default 30), --dry-run , --yes	Delete every scratch cache that's both closed and untouched for at least the given number of days. Reports what it removed — or, with <code>--dry-run</code> , what it <i>would</i> remove — and the total space reclaimed. Asks first (<code>--yes</code> skips; <code>--dry-run</code> never asks).
init	—	Write a starter <code>.swarmfile/cache.yml</code> into this mount's project root. Refuses to overwrite an existing file.

```
swarmfile scratch init
swarmfile scratch list
swarmfile scratch clear <name> --force
swarmfile scratch gc --older-than-days 14 --dry-run
```

Nothing here is unrecoverable project data: everything under a scratch cache is a re-buildable tool cache (a pnpm store, a cargo registry, ...), never project content, so clearing one only costs a slower next build, not lost work.

.swarmfile/cache.yml

A versioned, project-root config file — commit it alongside the project, like `.swarmfileignore` — declaring which tool caches `run` should redirect:

```
caches:
- name: pnpm
  env: PNPM_HOME
  subdir: pnpm
- name: cargo-registry
  env: CARGO_HOME
  subdir: cargo
```

Each entry needs a `name` (used only in messages), the `env` variable the tool reads for its cache location, and a `subdir` (relative, no `..`) under the shared scratch directory. A single bad entry — an unsafe subdir, a name collision with an earlier entry, or one of a handful of env vars this refuses to redirect on principle (`PATH` ,

`HOME` , `LD_LIBRARY_PATH` , and similar) — is dropped with a warning; it never takes the other, otherwise-valid entries down with it. Run `swarmfile scratch init` if you'd rather start from a commented example than write one by hand.

Only redirect READ-MOSTLY, CONTENT-ADDRESSED download caches here — a pnpm store, `~/ .cargo/registry` , a pip wheel cache. **Not** build `output` directories (`target/` , an in-progress `node_modules` install): sharing one of those across concurrent mounts risks lock contention, not just staleness.

changeset-idle

Changeset session-minting idle timeout — how long the engine waits after your last write before sealing the open changeset and starting a fresh one on the next write. This groups a burst of related saves into one commit rather than one commit per save. Only applies on a single-project mount. Distinct from `changesets list` (the commit history) — this is the settings command, not a log.

Subcommand	Flags	Description
<code>status</code>	—	This mount's current changeset idle timeout, in seconds.
<code>set</code>	<code><idle_secs: u64></code>	Set the idle timeout. <code>0</code> disables commit grouping (writes land ungrouped).

`changeset-idle set` restarts the engine to take effect (the changeset manager is only constructed at startup). Persists to the config file's `changeset_idle_secs` — see [Engine Config File](#).

attr-cache

How long (in seconds) the mount may trust cached file/directory attributes before revalidating — this controls how quickly a teammate's newly-saved file stops showing a stale directory listing. `0` means "leave the platform's own default alone" (not "disabled").

Subcommand	Flags	Description
<code>status</code>	—	This mount's current attribute-cache timeout, in seconds.
<code>set</code>	<code><attr_cache_secs: u64></code>	Set the timeout. <code>0</code> keeps the platform default.

`attr-cache set` restarts the engine to take effect, and applies only to macOS (`fuse-t`) / Linux (`vfs`) builds — it has no effect on a Windows/ `winfsp` mount. Persists to the config file's `attr_cache_secs` .

hydrate

Bulk-materialize a scope of this mount into the local cache, pinned against eviction until released — for a machine that needs to **read** that scope with no network, or that wants everything local before a job starts rather than fetching lazily during it. It does not reserve anything for writing: to also keep a scope editable with no hub, use [offline prepare](#) , which reserves the whole scope.

Subcommand	Flags	Description
<code>start</code>	<code>[path]</code> , <code>--yes</code> , <code>--wait</code> , <code>--sparse</code> <code><glob></code> , <code>--exclude</code> <code><glob></code>	Download and pin every file under <code>path</code> — relative to the mount, or a full path inside it (default: the whole project on the selected mount's branch — use the global <code>--mount <id></code> to target another mount). Prints a size estimate and a warning if it exceeds your cache cap, and asks for confirmation, unless <code>--yes</code> . <code>--sparse</code> / <code>--exclude</code> (the same globs as <code>mounts open</code>) narrow it to the matching paths; on a sparse mount the mount's own scope already applies, and a request can't widen it. Polls progress until done, except in <code>--json</code> mode, which returns as soon as the job starts unless <code>--wait</code> is given.
<code>status</code>	—	Progress of the current (or most recent) hydrate job.
<code>cancel</code>	—	Cancel a running hydrate job. A no-op if nothing is running.
<code>release</code>	<code>[path]</code>	Unpin everything hydrated under <code>path</code> on the selected mount's branch, making it evictable again. With no path, unpin everything this machine holds for the project — every branch, since the pin map is per machine and project-wide.
<code>free-space</code>	<code>[path]</code> , <code>--dry-run</code>	Free up space: delete the downloaded copy of everything under <code>path</code> so it's cloud-only on this computer, and stop keeping it offline. Files stay listed and download again when opened. Files with changes that haven't synced yet are skipped. With no path it covers the whole project and also clears older cached versions. <code>--dry-run</code> only reports what would be freed.

```
swarmfile hydrate start ./Projects/shots/seqA --yes
swarmfile hydrate start ./src --sparse "src/**" --exclude "src/fixtures/**"
swarmfile hydrate status
swarmfile hydrate release ./Projects/shots/seqA
swarmfile hydrate free-space ./Projects/shots/seqA --dry-run
```

Pinned blocks are never evicted by the cache's normal LRU sweep, so a large hydrate can push disk usage past your configured `cache cap` — `hydrate start` warns about this up front rather than silently growing past it. Pins are permanent until you run `hydrate release` , not time-limited.

sync

Saves that aren't reaching the cloud on their own — the cloud keeps refusing them, or their upload gave up — and a way to give one up.

Subcommand	Flags	Description
<code>stuck</code>	—	List stuck saves, one per file: the file's entry id, name and path, why it's stuck (<code>commit_refused</code> or <code>upload_failed</code>) and since when. A new file refused at a protected project's root is held, not lost: ask an owner to grant project-wide write, or create a folder to put it in — the full reason sentence is in <code>swarmfile status</code> and the Desktop App's Why did my save fail?
<code>discard</code>	<code><entryId></code> , <code>--yes</code>	Drop the file's stuck saves and put it back to its cloud version. The unsynced change on this computer is lost , so this asks first unless <code>--yes</code> .

```
swarmfile sync stuck
swarmfile sync discard 6f1c...e2 --yes
```

uploads

Every **live in-flight upload** on this mount — your own saves and other machines' — from the hub-synced pending pointers. This is the "what is still on its way" view that the Desktop App's **Uploading** badge summarizes, across everyone, not just this computer.

Flag	Description
<code>--wait</code>	Poll until nothing is uploading, instead of printing once. Useful before a checkout or a mirror in CI, where acting on a file that is still landing would read the old version.

```
swarmfile uploads
swarmfile uploads --wait # block until every in-flight upload has landed
```

`swarmfile uploads retract <path>` clears one file's pending pointer when the machine that was saving it is gone for good — the case where no upload will ever finish and nobody should keep waiting on it. It needs **write** access to the file, reads the live pointer first, and clears against that pointer's own manifest: if the upload has re-announced since, the retract is a no-op rather than an interruption, so it can never yank a genuinely live upload out from under its writer.

```
swarmfile uploads retract ./Projects/reel-04.r3d
```

A pointer nobody refreshes for 30 minutes stops blocking branch creates, merges from its branch, and copies even without this; retracting is for making that explicit and for the readers that wait on the pointer.

An upload that has **stalled** (a save this machine keeps failing to land) is a different list — see [sync stuck](#).

materialize

Write a ref's tree to a plain directory — **no mount, no FUSE**, so it works on a CI box with no kernel driver. The ref resolves through the same vocabulary as `show / restore (N / #N, HEAD~N, a branch, a tag, a changeset id, or a 64-hex commit hash; a prefix like hash: or tag: forces one kind)`, and it defaults to the selected mount's branch head.

A commit ref — a number, `HEAD` / `HEAD-N`, a tag, a hash, or another branch's name — writes **that commit's tree**, exactly as it was committed. Only the default (or the mount's own branch, e.g. `branch:main` on a main mount) writes the branch as the mount sees it right now, including saves that haven't been committed yet. A commit archived before it had a stored tree can't be materialized (`no_commit_tree`); pick a newer commit or a branch.

On a **protected** project, a member's checkout of a historical ref is filtered by the permissions that were in force at that commit — the hub replays its ACL log at the commit's own time, so files they could read then are written even if they're restricted now, and files they couldn't read then are left out. If the hub has no ACL history for that commit it prints a note that the listing was instead filtered against *current* permissions — never silently. Owners and service/CI callers (see [Runner](#)) are never filtered.

The destination must be new or empty, or a directory this tool already wrote: it keeps a `.swarmfile-materialize.json` marker and updates in place, pruning only paths it wrote before — so files your job created between runs (build outputs, caches) survive, and it never overwrites or deletes a directory it doesn't own.

A 64-hex hash is the most reproducible pin (tags can be repointed, branches move). Executable files (see `chmod`) are written executable on macOS and Linux. Exits non-zero if any file failed to materialize.

`--sparse <glob>` (repeatable) checks out only the paths that match — for a huge repo where a job needs one subtree. A path is written when it matches a glob or a directory in its ancestry does; `*.md` matches by name at any depth, `src/**` anchors to the root, and a leading `!` excludes a path an earlier glob included (gitignore-style). Everything else is neither fetched nor written, and a directory with no matching file isn't created. At most 64 globs, each at most 1024 characters, at least one non-`!`. `--exclude <glob>` is sugar for the `!` negations (with no `--sparse`, "everything except"). Omit either for a full checkout. The `clone --checkout` form takes both flags.

[illegible]

See [Runner \(Headless CI\)](#) for the full external-CI recipe, including the marker's ownership rules.

```
git status
```

How ready this project is to be served to git. Files at or above the git-LFS pointer threshold (default 1 MiB) would reach git as LFS pointers, and a pointer needs the file's SHA-256 object id. `git status` walks the selected mount's branch as the mount sees it and reports how many such files there are, how many still have no recorded id, and the first 50 of those with their sizes. It's the readiness check for turning on [git access](#), which refuses while any counted file is missing an id.

It reports the access state first, before the readiness counts: **enabled and frozen** (with the conversion version and pointer threshold), **not enabled** (with the command that turns it on), or **not available** on an end-to-end-encrypted project. When the hub can't be reached, that line is omitted rather than guessed at.

Flag	Meaning
<code>--threshold-bytes <N></code>	Count files of at least <code>N</code> bytes instead of 1 MiB.

Files saved through the desktop app (or imported with `swarmfile-migrate`) on a managed or unencrypted project get an id as they're saved. Files added through the web dashboard's uploader or produced by resolving a conflict get one too. Files with no recorded id show as missing until `swarmfile git backfill-oids` fills them

in. An end-to-end encrypted project never has ids — Swarmfile doesn't store a hash of E2E content — and says so; neither does content encrypted with a static key you supply yourself (`SWARMFILE_ENCRYPTION_KEY`).

It also reports how many commits on the branch are still **unconfirmed (unsigned)** — with how many of those are disputed and the oldest one's age — when there are any. A commit becomes servable to git as soon as it's derived, even before it's confirmed: a fetch returns the provisional hash and, from a second machine, confirms it. A teammate's app and an API key used in CI can confirm one the same way. A commit that stays unconfirmed for a long time is still worth a look; the warning never changes the exit code.

Exits `0` when every counted file has an id, `2` when some don't (or on an E2E project, or while the hub doesn't yet know the project's encryption tier). `--json` returns the counts and paths.

```
swarmfile git status
swarmfile git status --threshold-bytes 10485760 --json
```

git backfill-oids

Records the missing ids `git status` reports. For each file at or above the threshold that has no id, it reads the file through the desktop app's normal read path (downloading it if it isn't cached), computes its SHA-256, and sends the ids to Swarmfile. It works through the files in batches and prints progress as it goes. Project owners and admins only; an end-to-end encrypted project, or one encrypted with your own `SWARMFILE_ENCRYPTION_KEY`, never records ids and is refused.

Flag	Meaning
<code>--threshold-bytes <N></code>	Count files of at least <code>N</code> bytes instead of 1 MiB.

Exits `0` when nothing is left missing, `2` when some files couldn't be read (they're listed) or the project doesn't record ids. Re-running it is safe: files that already have an id are skipped.

```
swarmfile git backfill-oids
swarmfile git status # now reports every counted file as having an id
```

git enable

Turns on [git access](#) for the selected mount's project — the same one-way switch as the **Git access...** dialog in the web dashboard. Project owners and org admins only; needs a running engine with the project mounted. Enabling freezes the rules the project is converted to git with: the conversion version, the size above which a file becomes an LFS pointer, and where history starts. The record is immutable afterwards (changing any rule would change commit ids a clone has already seen), so the command prints exactly what was frozen and on which project.

Because the freeze cannot be undone, the command asks for confirmation first: `-y / --yes` skips the prompt, and a scripted run (`--json`, or no interactive terminal) must pass it rather than freeze silently.

Flag	Meaning
<code>--threshold-bytes <N></code>	Freeze the LFS pointer threshold at <code>N</code> bytes instead of 1 MiB.
<code>--no-backfill</code>	Refuse instead of recording missing object ids first. For scripts that want the fail-fast behavior; interactive runs don't need it.
<code>-y / --yes</code>	Skip the confirmation prompt. Required with <code>--json</code> and when stdin isn't a terminal.

It refuses, with the reason, while the project doesn't keep every saved version (turn that on from the project's history settings), when an older branch or tag head can no longer be read, when history is too long to check in one pass, and on an end-to-end-encrypted project — git access isn't available there.

If any at-or-above-threshold file still has no object id, the command records the missing ids itself first (the same work as `swarmfile git backfill-oids`: it hashes each file's content through the mount, fetching what isn't cached) and retries the enable once. `--no-backfill` skips that and surfaces the coverage refusal directly.

Exits `0` when the record was frozen, `2` on any refusal (the hub's sentence says which) or when the confirmation is missing, `1` when the hub couldn't be reached.

```
swarmfile git status # everything has an id? then:
swarmfile git enable
swarmfile git clone acme/website
```

git clone

Clone a project with git through the `swarmfile://` remote helper — the same as `git clone swarmfile://<org>/<project> [path]`, but it finds the `git-remote-swarmfile` helper next to this binary even when it isn't on your `PATH`. `<project>` is a project id, a bare project name (qualified with your configured organization), or `<org>/<project>`; `[path]` is the destination directory (git's own default when omitted). It runs `git clone` for you, so no running engine is needed. The project must have git access turned on — see [Cloning a Project with Git](#) for what a clone contains, how it signs in, how `git push` works from it, and what isn't supported.

```
swarmfile git clone acme/website
swarmfile git clone acme/website ./website
```

git import

Bring an existing git repository — GitHub, GitLab, a local path — into a new Swarmfile project, or into an existing one with `--into`. The project is created with git access already on and every version kept, and the source's default branch becomes the project's default branch (so upstream commit ids are preserved exactly, including a history whose default branch is `master`, `trunk` or anything else). The source is fetched into a temporary bare mirror with your own `git`, so the same

credentials a normal `git clone` would use — and a configured `git-lfs` — apply; each branch and tag is then pushed through the same `swarmfile://` helper `git clone` uses. It runs without a running engine, like `git clone`.

The repository name from the URL becomes the project name unless `--name` overrides it. Every branch the git view can hold is imported; one it cannot (submodules, Git-LFS, an octopus merge) is skipped and listed, while the same problem on the source's default branch refuses the import before anything is created — nothing can land without it. `--branches` narrows which branches are attempted (wildcards; the default branch is always imported) and `--no-tags` skips tags; `--submodules fail` / `--octopus fail` widen the refusal to side branches too. Annotated tags are imported as lightweight tags by default (the git view has no tag objects, so tag messages and signatures are not carried; the report says how many); `--annotated-tags skip` leaves them out and `--annotated-tags fail` stops the import before creating anything. Tags pointing at something other than a commit, or living only on a branch this import did not push, are skipped and listed. If anything fails to land, the temporary mirror is kept and the report names the retry: `swarmfile git import <url> --into <org>/<project>` resumes against what already landed; the project is resolved and its git access verified before the source is fetched, so a typo'd or disabled target is refused up front.

The import records the source it came from, so re-running the same command without `--name` / `--into` recognizes the project it made and syncs new upstream branches and commits into it instead of failing on the name. `--new` forces a fresh project anyway (a second project from one source), and an explicit `--name` always creates the named project.

The importer refuses, before creating anything, a source whose default branch uses Git LFS (LFS objects are not fetched yet), contains submodules, or has an octopus merge — each is a limit of git access itself. A project-name collision after the source is fetched is refused too; pass `--name`, or `--into` the existing project.

Flag	Description
<code>--name <NAME></code>	Project name (default: the repository name from the URL).
<code>--org <ORG></code>	Organization slug or id (default: your configured organization).
<code>--template <ID></code>	A creation template from the project-creation list (default: the <code>git</code> configuration without its starter <code>.gitignore</code> — an import needs the default branch to start empty, and the repository brings its own files, so a template that seeds files can't be imported into).
<code>--public</code>	Create a public plaintext project instead of the default private managed one (private projects need a plan with them).
<code>--into <ORG>/<PROJECT></code>	Push into an existing project instead of creating one; the project is resolved (before the source is fetched), so an unknown name is refused up front.
<code>--keep</code>	Keep the temporary mirror after a successful import.
<code>--annotated-tags <convert skip fail></code>	Annotated-tag handling: <code>convert</code> to lightweight (default), <code>skip</code> , or <code>fail</code> before creating anything.
<code>--new</code>	Create a fresh project even when this source was already imported (the default without <code>--name</code> / <code>--into</code> is to sync into the project this source made).
<code>--submodules <skip fail></code>	A branch containing submodules: <code>skip</code> (default) lists it and imports the rest, <code>fail</code> refuses the whole import.
<code>--octopus <skip fail></code>	Same for a branch containing an octopus merge.
<code>--branches <GLOB></code>	Import only branches matching the wildcard (<code>*</code> matches any characters, <code>/</code> included); repeatable, and the source's default branch is always imported.
<code>--no-tags</code>	Skip every tag (branches still import).

Exits `0` when every ref landed or was deliberately skipped (a tag pointing at a blob or tree is listed as skipped, not failed), `2` on a refusal (the sentence names the fix) or when a branch or tag push failed (the report names the `--into` retry), and `1` on a failure such as the source being unreachable. `--json` prints the structured report. The creation-only flags (`--name`, `--public`, `--template`) are refused alongside `--into`, which already names the project.

```
swarmfile git import https://github.com/carlini/js13k2019-yet-another-doom-clone.git
swarmfile git import git@github.com:acme/site.git --name website --org acme
swarmfile git import ./old-repo --into acme/website # retry a partial import
```

git url

Print a project's `swarmfile://` URL and exit, for wiring a remote by hand. `<project>` is a project id, a bare project name (qualified with your configured organization), or `<org>/<project>`. URL segments are printed ASCII-lowercased and percent-encoded where a character needs it (a display name with a space comes out as `%20`) — the canonical web-address form, which also repairs a caps-typed UUID (a full `swarmfile://` URL is passed through unchanged; the remote helper decodes an escaped segment before resolving it).

```
git remote add origin "$(swarmfile git url acme/website)"
```

git force-forward

The append-only form of `git push --force`, run inside a clone. A branch's tree can't be rewritten on Swarmfile, but the end state a force push wants — "the branch's tree is my tree" — can be landed as a *new* commit on top of the current head whose tree is your tip's, with your tip as its second parent. This command fetches the branch (which also snapshots uncommitted drive saves, like any fetch), builds that commit with git's own `commit-tree` (so its id is derived on this machine), and pushes it through the normal helper path: every commit stays reachable, teammates fast-forward instead of diverging, and your tip's history rides the `merged/<sha>` side branch the helper creates and archives. Only committed work is landed (a dirty worktree is ignored, like `git push`), and `--dry-run` still fetches — the dry run skips the push, not the fetch, so a fetch's usual effects apply.

Your local branch is deliberately not moved — the new commit's id cannot match your tip's — so the command prints the `git fetch <remote> && git reset --hard <remote>/<branch>` that aligns your clone. Pass a `BRANCH` (default: the checked-out branch), `--message <MESSAGE>` for the new commit's message (default `force-forward: <your tip's subject>`), `--remote <REMOTE>` (default `origin`), or `--dry-run` to run every check and push nothing. When your tip simply continues the remote head along its first parents (an ordinary fast-forward), a normal `git push` lands the same end state, and the command says so instead of building a graft; the diverged *merge* shape (merging the fetched branch into your line) does need the graft, and the push refusal for it names this command. When the branch already carries your tree there is nothing to do either, and the command just prints the realignment.

A clone can opt in to the same translation for a plain force push: `git config swarmfile.force-forward true`. With it on, `git push --force` (and `--force-with-lease`, whose lease check still runs client-side against the advertised head) pushes the tip as-is when it is an ordinary fast-forward, else builds the same forward commit; a note names the new id and asks you to `git fetch` and reset, since git's own bookkeeping points your remote-tracking ref at the tip you sent. Default off — force pushes stay refused.

Exits `0` when the push landed (or there was nothing to do), `2` on a refusal — detached `HEAD`, an unknown local branch or branch name, a fetch that couldn't find the branch, a push the helper refused — and `1` on a transport failure. A rewritten **root** commit cannot be landed this way: the replaced history shares no commit with Swarmfile, so the push refuses naming the unrelated root — push that as a new branch instead. `--json` prints the outcome — `pushed`, `fast-forward`, `already-at-tip` or `already-carries-tree` — with the commit ids; `pushed` and `already-carries-tree` carry `align`, the realignment command for scripts.

```
swarmfile git force-forward          # this clone's branch
swarmfile git force-forward release -m "keep our tree"
swarmfile git force-forward --dry-run
git config swarmfile.force-forward true      # let plain --force use the same path
```

git index

Derives commits' Git-ready trees on this machine and hands them to Swarmfile, so a committed version can be served to git directly. For each commit it reads the branch state as of that commit's version, builds the tree objects (directories, symlinks and the large-file pointers, splitting very large directories into segments), uploads anything the hub doesn't already have, and claims the result. A second machine — a teammate's app, a fetch, or an API key used in CI — confirms the same value.

With no argument it works on the selected mount's branch head and every unconfirmed commit behind it (through merges too), back to the nearest commit that is already confirmed or that you claimed yourself. Each commit is derived independently and the claims go oldest first, so a chain someone else pushed is confirmed in one run. A run handles up to 1,000 commits and prints one line per commit. Given a commit's sequence number (`#42` or a bare number), it handles just that commit.

The git remote helper also does this by itself for a small unclaimed chain: when a git clone, fetch, push or `ls-remote` meets a branch head that has never been derived, it derives and claims up to 200 commits inline (a clone that set `swarmfile.confirm=false` opts out). This command is for longer chains, or a machine without the Desktop App.

Requires a running engine and a managed or unencrypted project. Exits `0` when every claim was accepted (including ones still waiting for a second machine to confirm), `2` when the hub disputed or rejected one (later commits in the run aren't claimed), and `3` when the hub couldn't be reached.

A commit printed as `rejected` (`resolved_against` in `--json`) was already settled to a different value when the hub resolved a dispute over it, so this machine's claim is kept only as evidence; compare the two with `git disputes SEQ --all` and report both hashes, since it is a client bug unless the project's history changed. If the settlement left the commit provisional, an owner who believes it is wrong can override it with `git clear-claim`; a settlement that confirmed the commit is final.

A commit printed as `provisional` (your earlier claim was kept) (code `self_disagreement` in `--json`) means this machine derived a different value than its **own** earlier claim for the same commit. One identity deriving two values for one commit is a client bug, not a dispute: the hub keeps the first claim, records the disagreement as evidence (readable with `git disputes SEQ --all`), and does not alert owners. Report it with both hashes.

```
swarmfile git index                # the branch head and every unconfirmed commit behind it
swarmfile git index '#42'          # one commit
swarmfile git index --json
swarmfile git index --verify-full  # ignore the cache; re-derive and check every object
```

`--verify-full` re-derives without the incremental/sparse caches and checks each derived object against its content address before uploading — a slow but complete self-check for a suspicious or freshly upgraded machine.

git clear-claim

Owner/admin only: resets the commit claim at a sequence number on the selected mount's project. Most claims are never cleared by hand: Swarmfile's server (for a pushed commit) or a teammate's fetch confirms a provisional one, and a disputed one is settled by Swarmfile deriving it again on the server, usually within a minute or two. This is the escape hatch for a claimed commit that can't be served: a **disputed** row (two derivations disagree), a claim whose root tree is gone, or a **stranded snapshot-mint head** — `git push` mints the branch's live tree as a snapshot when it has drive saves, and if the push is then refused, that snapshot stays as the provisional head; clearing it also rolls the branch head back to the snapshot's parent, so the next fetch mints afresh from the live tree. A healthy claim — including a member's own — is never discarded: the hub answers `cleared: false` and says why.

Clearing resets the hub's record of the claim, so it is confirmation-gated: `-y / --yes` skips the prompt, and a scripted run (`--json`, or no interactive terminal) must pass it. Dispute evidence is kept, not deleted: an owner or admin can read both values, their trees, parents and claimants before or after a clear with `git disputes` (add `--all` once it is cleared), or in the **Commit disputes** section of the project's Commits page on the web.

Argument/flag	Meaning
<code>SEQ</code>	The changeset sequence number to clear.
<code>-y / --yes</code>	Skip the confirmation prompt. Required with <code>--json</code> and when stdin isn't a terminal.

Exits `0` when the claim was cleared (an already-confirmed or archived seq prints that instead), `2` on a refusal (not an owner/admin) or when the confirmation is missing, `1` when the hub couldn't be reached.

```
swarmfile git index '#42'          # see what the hub holds first
swarmfile git clear-claim 42       # reset it; a stranded snapshot head rolls back
swarmfile git clear-claim 42 -y
```

git disputes

Owner/admin only: lists the commit disputes on the selected mount's project, newest first. A dispute means two derivations of the same commit produced different hashes, so git access can't serve that commit until Swarmfile settles the dispute (usually within a minute or two, by deriving the commit again on the server) or an owner or admin clears its claim. Each block shows the commit's sequence number, when the dispute was recorded, the claim's current status, and both sides: the challenging claim and the one it contradicted, each with its commit hash, tree, parents and claimant. A claimant is a user, an API key, or Swarmfile's server-side derivation.

By default only open disputes are listed: not yet cleared, on a commit that is still disputed or provisional. Resolve one Swarmfile could not settle, or override a settlement that left the commit provisional, with `git clear-claim`; the evidence is kept, and Swarmfile derives the commit again on its own (unless its own derivation was part of the dispute), as does the next `swarmfile git index` or a teammate's fetch.

Argument/flag	Meaning
<code>SEQ</code>	Only the disputes for this commit sequence number.
<code>--all</code>	Include cleared and evidence-only disputes.

`--json` prints the hub's list unchanged. Exits `0` on success (an empty list included), `2` on a refusal (not an owner/admin), `1` when the hub couldn't be reached.

```
swarmfile git disputes          # open disputes on this project
swarmfile git disputes 42 --all # every record for commit 42, cleared ones included
swarmfile git disputes --json
```

offline

Pack & Go: materialize a scope **and** reserve the whole scope with one long-lived offline reservation, so it stays editable with the hub unreachable. `hydrate` above only makes content readable offline; this is the form that also makes it writable. It runs as a background job, and the scope follows the selected mount's branch (target another mount with the global `--mount <id>`).

Subcommand	Flags	Description
<code>prepare</code>	<code>[path]</code> , <code>--duration-secs <N></code> , <code>--offline-days <D></code> , <code>--no-wait</code>	Materialize and reserve. The lease defaults to 7 days, clamped at 30; a lease is only renewed while online, so <code>--offline-days</code> , when longer than the lease, extends the lease to cover the trip (an explicit shorter <code>--duration-secs</code> is never overridden — it warns instead). Prints how many files the reservation covers. If another machine holds an overlapping reservation or is editing a file inside the scope, the prepare is refused naming the holder and nothing is reserved. Then polls to completion unless <code>--no-wait</code> .
<code>status</code>	—	Progress of the current (or last) job — <code>interruptedScope</code> names one a restart interrupted — plus what this engine holds right now (file edit locks and folder/project reservations, with an expiring-soon and renewal-failure warning).
<code>cancel</code>	—	Ask a running prepare to stop — reservations already taken are kept, and the content download this prepare started is stopped too (only that job, never a newer one).
<code>resume</code>	—	Complete a prepare a restart interrupted; re-reserves the scope idempotently.
<code>return</code>	<code>[path]</code>	Release reservations: at or under one path, or with no path every reservation this engine holds (per-file locks and scope reservations alike). In-scope saves are given a few seconds to land first; any still syncing are reported.

```
swarmfile offline prepare ./Projects/shots/seqA --offline-days 14
swarmfile offline status
swarmfile offline return ./Projects/shots/seqA
```

fetch

Pull a byte range of one file into the local cache — including from a version somebody else is **still uploading**.

On a slow link a large upload takes hours or days, and until it finishes the new version isn't readable at all. This asks the uploading machine to send the part you need *first*, then downloads it as it lands. What you get back is the time it takes to send that range, not the time it takes to send the file.

Flag	Description
<code>--tail <bytes></code>	Fetch the last N bytes. The common case for media, where the part you need is at the end. Resolved against the <i>in-flight</i> size, which is why it's a flag rather than arithmetic you do yourself — see below.
<code>--start <bytes></code>	First byte of the range. Default <code>0</code> . Conflicts with <code>--tail</code> .
<code>--len <bytes></code>	How many bytes from <code>--start</code> . Clamped to the end of the file. Conflicts with <code>--tail</code> .
<code>--version <which></code>	<code>newest</code> (default — the in-flight version if there is one, otherwise the committed one), <code>pending</code> (fail unless something is uploading), or <code>head</code> (the committed version only).
<code>--follow</code>	Track a moving reader instead of fetching a fixed range: keep a buffer ahead of wherever it is and re-target when it seeks. <code>--start</code> becomes the initial position; <code>--len</code> and <code>--tail</code> are rejected together with <code>--follow</code> .

```
# The last 200 MB of a file a colleague is still uploading
swarmfile fetch ./Projects/reel-04.r3d --tail 209715200

# A specific window of the committed version
swarmfile fetch ./Projects/site.rvt --start 0 --len 52428800 --version head
```

The command polls until every block has landed, printing progress as it goes:

```
$ swarmfile fetch ./Projects/reel-04.r3d --tail 209715200
19/160 blocks, 25.8 MiB fetched (waiting on the upload)
...
fetch complete: 160/160 blocks, 200.0 MiB now local
```

"waiting on the upload" is the normal state, not a stall. The blocks you asked for may not exist yet — the other machine is still sending them at its uplink speed — so this command can legitimately run for a long time. It gives up only when nothing is arriving *and* the hub reports that the upload has stopped, and says so rather than hanging.

Why `--tail` rather than working out the offset yourself: the size to measure back from is the **in-flight** version's, and the file's committed size is still the old version's. A number you compute locally would measure from the wrong end of the wrong file.

`--version pending` fails outright when nothing is uploading, rather than quietly giving you the committed version. That is deliberate: somebody who asked for the version in progress and silently received yesterday's bytes has no way to tell.

Fetched blocks are pinned exactly like `hydrate`, so they aren't evicted by the cache's LRU sweep before you open the file — and they are released the same way, with `swarmfile hydrate release <path>`. One job runs at a time; starting a second while one is running is refused rather than queued.

You can also do this without the CLI: reading the file normally through the mount always works, and gets the committed version. `fetch` is for the case where the version you want is still on its way. (A mount also streams a file that has never finished uploading — see [the config reference](#).)

Following a moving reader

A fixed range is the right shape for "give me the last 200 MB". It's the wrong shape for watching a file as it uploads, because a viewer plays forward and seeks. `--follow` keeps a buffer ahead of a read position you report as it moves, and re-targets when it jumps.

Command	Flags	Description
<code>fetch <path> --follow</code>	<code>--start <bytes></code> , <code>--version <which></code>	Start following from <code>--start</code> (default <code>0</code>). Returns once the job is running; it does not complete on its own.
<code>fetch-seek</code>	<code><position: u64></code>	Report where the reader has got to. The buffer re-targets there immediately. A no-op if nothing is following.
<code>fetch-status</code>	—	Progress of the current (or last) <code>fetch</code> . The way to watch a <code>--follow</code> job.

```
swarmfile fetch ./Projects/reel-04.r3d --follow --version pending
swarmfile fetch-seek 5000000
swarmfile fetch-status
```

For a follow job, `fetch-status` reports `position`, `bufferedBytes` and `readRateBps` rather than a block count. **bufferedBytes is the number worth watching:** it counts contiguous bytes from the read position, so it answers "will this play" — where a block total does not. A hole two blocks ahead stalls a reader however full the rest of the window is.

Buffer depth is derived from how fast the reader is actually consuming, not from a fixed byte count: 30 seconds of a 6 Mbit/s proxy and 30 seconds of a 200 Mbit/s master are very different numbers of bytes, and a single figure would be useless for one of them. Set the seconds with `SWARMFILE_STREAM_BUFFER_SECS` (default 30). A seek is not counted as playback, so scrubbing doesn't inflate the estimate.

A follow job ends when you cancel it, when the upload finishes and the buffer reaches the end of the file, or after five minutes with no `fetch-seek` — a paused viewer reports nothing, and one job runs at a time, so an abandoned follow would otherwise block every later fetch.

seed

NAS/seed-node mode.

Subcommand	Description
<code>status</code>	Current seed-mode state.
<code>enable</code>	Turn on seed mode. Requires the Pro plan or above, and is mutually exclusive with Cloud-only mode.
<code>disable</code>	Turn off seed mode.

`seed enable / disable` restart the engine.

A headless cache is easiest to set up from the dashboard: **Settings → Office caches** generates the complete `seed.env` (scope and credential included) and lists the fleet's liveness and served bytes. `seed enable` is the toggle for an engine whose scope is already configured.

watch / unwatch / watches

Command	Description
<code>watch <path></code>	Watch a file or folder.
<code>unwatch <path></code>	Stop watching a file or folder.
<code>watches</code>	Everything you're currently watching in this project.

lock / unlock / locks

An explicit, up-to-30-day edit lock — like `git lfs lock` / Perforce's `p4 edit`. Not the same as the short-lived, automatic lock a plain write takes (from an app's first write to the file until it closes it); this one you take and release yourself.

Command	Flags	Description
<code>lock <path></code>	<code>--duration-secs <N></code>	Reserve exclusive edit rights on a file, even offline. Defaults to 7 days, clamped to a 30-day maximum. Refused (exit 2) if someone else already holds it. Re-running <code>lock</code> on a file you already hold extends it, and a held lock is auto-renewed every 6 hours while the engine is online.
<code>unlock <path></code>	—	Release a lock taken with <code>lock</code> .
<code>locks</code>	—	Every file currently locked — for editing, or being written somewhere — in this project's current branch.

```
swarmfile lock ./Projects/tower-a.rvt
swarmfile lock ./Projects/tower-a.rvt --duration-secs 86400 # 1 day
swarmfile locks
swarmfile unlock ./Projects/tower-a.rvt
```

If someone else needs the file back before you release it, they can ask via [unlock-request](#) below — there's no force-release from another user's side.

chmod

Mark a file executable or not — the same bit `chmod +x` sets on the mounted drive, and the one git records as mode `100755` (a script, a build tool) versus `100644`. Swarmfile stores only this one bit: there are no separate read/write bits per user.

Command	Description
<code>chmod +x <path></code>	Make a file executable.
<code>chmod -x <path></code>	Make it a plain, non-executable file.

```
swarmfile chmod +x ./tools/build.sh
swarmfile chmod -x ./docs/notes.txt --json
```

- Files only: a folder or symlink is refused (`not_a_file`).
- Works offline. The change shows at once on this machine and is sent when the hub is reachable again; the summary line (or `"queued": true` under `--json`) says when that's still pending.
- On the macOS and Linux drive, `chmod` does the same thing: any execute bit (`u+x`, `g+x` or `o+x`) makes the file executable, and none makes it not. The drive shows executable files as `rw-r--r--x` and every other file as `rw-r--r--`.
- Windows has no execute bit, so `swarmfile chmod` is how you set it there. Editing a file on Windows keeps its bit, and so does an app that saves through a temporary file and renames it over the original.

comment / comments

Command	Flags	Description
<code>comment <path></code>	<code>-m</code> , <code>--message <String></code> (required)	Post a comment on a file.
<code>comments <path></code>	<code>--limit <N></code> , <code>--offset <N></code>	List comments on a file, newest first. The hub caps one page at 199; <code>--offset</code> pages through the rest.

trash

Soft-deleted entries in this mount's project — the headless equivalent of the Desktop App's Trash panel.

Subcommand	Flags	Description
<code>list</code>	<code>--limit <N></code> (default 50), <code>--offset <N></code>	Your own soft-deleted entries, newest-deleted first.
<code>restore <entry-id></code>	<code>--with-nearby</code>	Undelete back to its original location. Refused with <code>name_conflict</code> / <code>parent_deleted</code> (exit 2) if something else has since taken that name or the parent folder is itself gone — either way the response names the conflicting entry so you can resolve it by hand. A folder deleted from the desktop drive is deleted file by file, so restoring it brings back the folder and lists the items deleted inside it just before it (<code>nearbyDeleted</code> under <code>--json</code>); <code>--with-nearby</code> restores those too.
<code>purge <entry-id></code>	<code>--preview</code> , <code>--yes</code>	Permanently delete ("Delete Forever") — not recoverable, unlike a normal delete. <code>--preview</code> reports how many descendants a folder would take with it, without removing anything. Asks first; <code>--yes</code> skips the prompt (or refuses on a non-interactive stdin / under <code>--json</code>).

```
swarmfile trash list --limit 100
swarmfile trash restore 3f9a2b --with-nearby
swarmfile trash purge 3f9a2b --preview
swarmfile trash purge 3f9a2b --yes
```

unlock-request

Ask someone to release a file they have locked — distinct from `lock / unlock` above (edit locks *you* hold); this is the "please give it back" workflow for a lock someone **else** holds, the same feature the Desktop App's request-unlock modal drives.

Subcommand	Flags	Description
<code>create <path> [-m, --message]</code>	—	Ask the current holder to release <code>path</code> . If it just freed up on its own, the hub reports that directly (<code>noLock: true</code>) instead of creating a request nobody needs to answer.
<code>list [--status] [--entry-id] [--limit]</code>	<code>--status</code> (default <code>pending</code>)	Unlock requests you're involved in, as requester or holder.
<code>respond <entry-id> <request-id> --decision <granted denied></code>	—	As the current holder , grant or deny a pending request against a file you have locked. Exit 2 if it's already resolved or isn't yours.
<code>cancel <request-id></code>	—	As the requester , withdraw a still-pending request. Exit 2 if it's already resolved.

```
swarmfile unlock-request create ./Projects/tower-a.rvt -m "need to push a fix before EOD"
swarmfile unlock-request list --status pending
swarmfile unlock-request respond 3f9a2b req_88 --decision granted
```

presence (alias who)

Who's currently editing what in this project.

check-ignore <path>

Report whether a path is excluded from sync by `.gitignore / .swarmfileignore` , and by which rule — the diagnostic counterpart to `ignore-clean` . Answers "why isn't this file syncing?" the way `git check-ignore -v` does.

```
swarmfile check-ignore ./Projects/node_modules/pkg.js
```

ignore-clean [path]

A `git rm --cached` equivalent: remove already-hub-synced content that matches this machine's own `.gitignore / .swarmfileignore` sync-exclusion rules (see `check-ignore`) — for content that synced before the matching rule applied, or from a machine whose ignore rules differ. The hub delete is a **soft delete** (recoverable from Trash for a while), and a matched directory's whole subtree is removed in one call rather than entry by entry. Prints a size estimate and asks for confirmation first, unless `--yes` ; without `--yes` a non-interactive stdin or `--json` refuses (exit 2), and `--dry-run` prints the estimate without confirming or deleting.

Flag	Description
<code>[path]</code>	Mount-relative path to narrow the scope to a folder or file. Omit for the whole project.
<code>--dry-run</code>	Print the size estimate and exit without deleting anything — the preview of exactly what <code>--yes</code> would remove.
<code>--yes</code>	Skip the confirmation prompt.
<code>--force</code>	Cross-machine safety override. Alongside <code>--yes</code> , proceed even if a top-level candidate was last modified by someone other than you within the last 7 days. Without <code>--force</code> , <code>--yes</code> refuses in that case. Interactive mode is never gated by this — a human already sees the same who/when detail and answers the prompt regardless.

```
swarmfile ignore-clean
swarmfile ignore-clean ./Projects/old-import --yes
```

completions <shell>

Generate a shell completion script. Works without a running engine. Supports `bash` , `zsh` , `fish` , `elvish` , `powershell` .

```
swarmfile completions zsh > ~/.zsh/completions/_swarmfile
```

doctor

Connectivity diagnostics — asks the running engine to run its probe suite. For a machine where the engine isn't running, use the standalone `swarmfile-doctor` binary instead.

log

The commit log as a compact, colored feed (like `git log --oneline`), scoped to the branch this mount is on. Each line shows the commit's 12-char hash (the full hash is what `hash: refs` and CI pins use).

Argument / Flag	Description
<code>[REF]</code>	Start the log at a commit and walk back from it. Full ref vocabulary (<code>#N / HEAD~N</code> , a branch, a tag, a changeset id, or a 64-hex hash). A branch or tag lists that branch's history (<code>log feature</code> on a main mount shows feature's commits); a bare number, hash, or changeset id walks the mount's branch. <code>--branch</code> wins when given.
<code>--limit <N></code>	Default <code>20</code>
<code>--branch <name></code>	Log another branch's history (read-only) instead of the mount's.

tail

Watch this project's activity live — comments, watched-file changes, lock conflicts. Ctrl+C to stop.

This is the engine's **local** event bus for the mounted project, not the org-wide Activity feed — for that, use `activity` (`list / export / tail`). The global `--mount` flag has no effect on `tail` : the stream covers every open mount's activity, not just the current one.

Headless engines

api-key

Project-scoped API keys, for an engine that runs where nobody can complete a browser sign-in — a render-farm node, a CI runner, a build box. A machine holding one authenticates without an interactive OIDC session at all.

Subcommand	Description
<code>list</code>	Every key for this mount's org. Never shows key material.
<code>create <name> [--project <id>]</code>	Mint a key. <code>--project</code> defaults to this mount's own project; it takes the project id , not its slug (an id the org doesn't have is refused with <code>project_not_found</code>).
<code>revoke <id></code>	Revoke by id (from <code>list</code>). Immediate — the key stops authenticating right away.
<code>rate-limit <id> [max] [--clear]</code>	Set (or, with <code>--clear</code> , clear) that one key's own request ceiling, in requests per minute. Capped at the plan's per-user maximum — use it to let a render-farm/agent key burst above the plan's flat default.

Minting and revoking require an **admin or owner** role in this mount's org. The same mint/list/revoke lives in the dashboard, under API keys.

Minting is capped by the plan's **headless-key allowance** — Free 1, Starter 2 per seat (min 5), Pro 5 per seat (min 10), plus any committed key packs — and a mint past the cap is refused with `keys_cap` ; revoke an unused key or add a pack under Settings → Billing. Revoking frees the slot immediately, and a downgrade never revokes keys you already have — it only blocks minting new ones past the new allowance.

`create` prints the raw secret **once**. The hub never returns it again — copy it before you scroll. With `--json` the secret is in the response body's `key` field, which is what makes it scriptable; treat that output as the credential it is.

A key is walled off to exactly one project. It cannot read or write anything else in the org — not even projects the admin who minted it can see — so a compromised farm node is confined to the project it was built for, and revoking a key stops that one machine without touching anybody's own session.

Using one. Set it as `SWARMFILE_API_KEY` on the machine, and the engine skips sign-in entirely: no browser, no refresh token, no impersonating whoever happened to mint it. If both `SWARMFILE_API_KEY` and `SWARMFILE_OIDC_REFRESH_TOKEN` are set the API key wins, and the engine says so in its log rather than choosing quietly — a leftover personal refresh token on a shared machine is the more likely mistake of the two.

```
swarmfile api-key create farm-node-12 --project proj_9f2a
```

If the mount covers more than one project, `--project` is required: there is no single obvious answer, and guessing one would scope a credential to the wrong place.

See [Deployment Topologies](#) for how this fits with per-node config files and seed nodes.

project

Project lifecycle and settings in this mount's org — the headless/scripted equivalent of the Desktop App's project settings.

Subcommand	Flags	Description
create <name> [--e2e] [--public] [--commit-mode] [--project-type] [--keep-full-history] [--description] [--git-compatible] [--acl-mode] [--hub-only] [--residency] [--protect-default-branch] [--initial-branch] [--option]	--e2e; --public; --commit-mode <auto-commit staged>; --project-type <revit_worksharing>; --keep-full-history <true false>; --description <text>; --git-compatible; --acl-mode <open protected> (Open by default; protected needs acls); --hub-only <true false>; --residency <eu us fedramp fedramp-high> (pins this project's region at creation — per-project and independent of the org's, fixed thereafter; needs the plan's residency capability <i>and</i> provisioned storage for the region); --protect-default-branch; --initial-branch <name> (the default branch's name — most projects keep <code>main</code> ; fixed at creation); --option <key=value> (repeatable — any create option by manifest key, e.g. <code>template=blank</code>)	media
list	—	Every project in this mount's org, with the currently-active one marked.
get <project-id>	—	Full settings for one project by id (see <code>list</code> for ids).
rename [name] [--slug]	—	Rename this mount's project and/or its slug.
archive / unarchive	—	Soft-archive this mount's project, or reverse it.
delete [--yes]	--yes	Permanently delete this mount's project and all its files, for everyone in the org. A hub-side soft delete: the project leaves active use immediately and its files, history, git-LFS objects, and stored data are permanently purged 30 days later. Recovery isn't self-serve — contact support before the purge window closes to discuss it (it isn't guaranteed; the 30-day window is fixed, and a shorter per-user trash setting does not shorten it). Prints a warning and asks for confirmation first, unless <code>--yes</code> ; refuses outright in <code>--json</code> mode without <code>--yes</code> rather than reading stdin. Consider <code>archive</code> instead if you just want to hide it.
set-keep-full-history <true false>	—	Opt this project's history out of retention GC (<code>true</code> , every past version stays revertable forever) or back into it (<code>false</code>). New projects default per the organization's new-project default (keep every saved version unless an owner changed it); create <code>--keep-full-history false</code> is the create-time opt-out, and this flips an existing project either way.
set-type <revit_worksharing none>	—	Mark this project as Revit-worksharing (native OS-file-lock enforcement, hub-cross-machine), or clear it with <code>none</code> .
set-commit-mode <mode>	—	Lock this project to a commit mode — a branch lock overrides this while set. <code>mode</code> is <code>auto-commit</code> , <code>staged</code> , or omitted/ <code>unset</code> to clear.
show [project-id]	--json	Print this project's effective settings and where each value came from — the CLI view of the Desktop App's project settings. Defaults to the active project; <code>--json</code> for scripting.
set <key=value>...	—	Set one or more settings in a single call. Supported keys: <code>commit-mode</code> , <code>keep-full-history</code> , <code>type</code> , <code>name</code> , <code>slug</code> , <code>description</code> , <code>archived</code> , <code>hub-only</code> (<code>true</code> / <code>false</code> / <code>inherit</code> ; <code>false</code> and <code>inherit</code> both mean "inherit the org default" — a project can only add hub-only, never force peer-to-peer over an org hub-only policy), <code>allow-external-collaborators</code> (<code>true</code> / <code>false</code> / <code>inherit</code> ; the per-project guest policy, where <code>inherit</code> follows the org's), <code>allow-public-rfi</code> and <code>allow-public-access-requests</code> (<code>true</code> / <code>false</code> ; the public project's RFI and access-request toggles, on by default). Refuses an unknown key rather than silently ignoring it.
rotate-key	—	Rotate this mount's E2E project key : mint a new generation and re-wrap it to every current device plus the org recovery key. Owner/admin and E2E only. Content written before the rotation stays sealed under the old generation and is not re-encrypted, so a device that already unwrapped the old key can still read it.

`rename` / `archive` / `unarchive` / `delete` / `set` / `set-keep-full-history` / `set-type` / `set-commit-mode` / `rotate-key` are owner-or-org-owner-gated: exit 2 on `forbidden` or `not_found` , exit 1 for anything else. `rotate-key` also exits 2 on a 409 (`stale_generation` — another device rotated first — or a recipient set that kept changing).

```

swarmfile project create "Tower Block A"
swarmfile project create "Secret Project" --e2e
swarmfile project create "Public Dataset" --public
swarmfile project create "Lean Project" --keep-full-history false
swarmfile project create "Model" --project-type revit_worksharing --description "Central model"
swarmfile project create "Repo" --git-compatible --acl-mode protected --protect-default-branch
swarmfile project create "EU Dataset" --residency eu --option hub-only=true
swarmfile project create "Repo" --git-compatible
swarmfile project show
swarmfile project set commit-mode=staged keep-full-history=false hub-only=inherit
swarmfile project rename "Tower Block A - Phase 2"
swarmfile project set-keep-full-history true
swarmfile project rotate-key

```

config

YOUR personal default commit mode for this mount's project — advisory only. Distinct from `project set-commit-mode / branch set-commit-mode` above, which are admin **locks**: those, when set, always win over this personal preference. Also distinct from `changeset-idle / attr-cache /etc.` above, which are local, per-machine engine-behavior knobs, not a hub-side preference tied to your account.

Subcommand	Flags	Description
<code>set-commit-mode [mode]</code>	—	Set (or, omitted/ unset , clear) your personal default. <code>mode</code> is <code>auto-commit</code> or <code>staged</code> . Refused (exit 2, <code>policy_enforced</code>) if an admin lock already resolves for this project+branch — a preference is never silently set unable to take effect.
<code>get-commit-mode</code>	—	Shows your personal preference (<code>pref</code> , null when unset) and the effective mode with its source: the response's <code>effectiveMode</code> is what this mount's next save or <code>changelist open</code> will use, and <code>effectiveSource</code> is <code>branch</code> or <code>project</code> for an admin lock (a lock always wins), <code>user_pref</code> when your own <code>set-commit-mode</code> default applies, or <code>unenforced</code> when nothing is set.

```

swarmfile config set-commit-mode staged
swarmfile config get-commit-mode
swarmfile config set-commit-mode unset

```

runner

Read-only run history for the headless runner-as-CI daemon — see [Runner \(Headless CI\)](#) for the full walkthrough. Jobs execute unsandboxed as the daemon's own local user (see [Security](#)). There is no `runner create / runner edit` here: rules live in `.swarmfile/runner.yml` on the watched branch, not on the hub, so there's nothing on the hub side for this command to create or edit.

Subcommand	Flags	Description
<code>runs</code>	<code>--branch <name></code>	Run history — every branch by default, or one branch with <code>--branch</code> . Each row carries the matched rule's name, the triggering commit, status (<code>running</code> / <code>success</code> / <code>failure</code> / <code>timed_out</code>), exit code, and a truncated log.

```

swarmfile runner runs
swarmfile runner runs --branch main

```

ci

Hosted CI (`.swarmfile/ci.yml` workflows that run on Swarmfile's own containers — see the [Hosted CI guide](#)). Requires the engine to be signed in: `reads`, `settings / size list`, `run` and `cancel` need a signed-in member; `enable/disable`, `secrets`, `variables` and `size` **writes** are owner/admin on the hub, and hosted CI stays off for a project until an owner enables it. In some organizations hosted CI is unavailable entirely, in which case these commands answer `404 ci_not_available` . The workflow dialect is documented in the [Hosted CI guide](#).

Subcommand	Flags	Description
list	<code>--branch <name></code> , <code>--limit <n></code>	Run history, newest first — every branch by default. Rows carry the workflow name, run number, commit, status and a blocked reason when admission refused.
status	<code><run-id></code>	One run plus its jobs: status, size, exit code, billed seconds and metered charge.
cancel	<code><run-id></code>	Ask the hub to cancel a queued/running run.
rerun	<code><run-id></code>	Re-run a finished run as a NEW run at the same branch, commit and changeset — every job starts fresh. The hub re-runs the trigger gates, so an E2E project or an unprotected branch without the opt-in refuses exactly like the original trigger (with the same typed code).
retry	<code><job-id></code>	Retry ONE failed job in place (owner/admin). The job's attempt count bumps, its per-attempt fields reset, and the run reopens and re-finalizes when the job finishes. <code>success</code> / <code>running</code> jobs cannot be retried.
run	<code>--branch <name></code> , <code>--input NAME=VALUE</code>	Manually dispatch the branch head's workflow (<code>workflow_dispatch</code>). <code>--input</code> is repeatable for workflows declaring <code>on.manual.inputs</code> ; the hub coerces values to the declared type.
logs	<code><job-id></code> , <code>--offset <bytes></code> , <code>--follow</code>	One job's log chunk on stdout (pipe-friendly); <code>--offset</code> continues where the last chunk ended. <code>--follow</code> keeps reading from the offset cursor (~2 s apart) and stops once the job is terminal and the tail is drained; Ctrl-C stops cleanly.
settings	—	The project's hosted-CI settings.
enable	<code>--allow-unprotected</code>	Enable hosted CI for the project. <code>--allow-unprotected</code> is the loud opt-in for branches without effective protection (workflows can reference secrets, so leave it off unless you mean it).
disable	—	Disable hosted CI for the project.
secret list	—	Secret names and metadata only — values are write-only and never printed.
secret set	<code><name></code> , <code>--from-env <VAR></code>	Store a secret value: a hidden prompt when run in a terminal, <code>--from-env <VAR></code> to read it from that environment variable, or piped stdin for scripts. The value never echoes or lands in shell history.
secret delete	<code><name></code>	Remove a secret.
variable list	—	Names and values of the non-secret config.
variable set	<code><name> <value></code>	Create or update a variable.
variable delete	<code><name></code>	Remove a variable.
size list	—	Org-defined custom runner sizes.
size set	<code><name></code> , <code>--vcpu</code> , <code>--memory</code> , <code>--disk</code>	Create or update a custom size (quarter-vCPU steps; the platform enforces the ceilings).
size delete	<code><name></code>	Remove a custom size.

```

swarmfile ci list --branch main
swarmfile ci status 018f3c2e-...
swarmfile ci logs 018f3c2e-... --offset 65536 | less
swarmfile ci logs 018f3c2e-... --follow           # live until the job ends
swarmfile ci rerun 018f3c2e-...                   # clone a finished run
swarmfile ci retry 018f3c2e-...                   # re-queue one failed job
swarmfile ci run --branch main
swarmfile ci secret set NPM_TOKEN                 # hidden prompt for the value
swarmfile ci secret set NPM_TOKEN --from-env NPM_TOKEN # read it from the environment
swarmfile ci variable set REGION eu
swarmfile ci size set big --vcpu 4 --memory 12 --disk 20

```

office

Peer-discovery office grouping — applies immediately, no restart.

Subcommand	Flags	Description
status	—	This mount's current <code>office_id</code> and whether <code>lan_from_office</code> is on.
set	<code><office_id></code> , <code>--lan-from-office</code>	Set the office-grouping label this mount announces for peer discovery. Pair with <code>--lan-from-office</code> on a network that blocks mDNS multicast — see <code>ec-placement</code> above.

Seed nodes are the one exception to `--lan-from-office`: a seed node joins another peer's placement roster purely by matching `office_id`, with no flag and no shared network required at all — see [Self-Hosted Seed Nodes](#) for using this to give a cloud-hosted render farm deliberate redundancy among its own nodes.

acl

Folder-ACL project mode and the ACL change log — the headless equivalent of the dashboard's ACL settings panel.

Subcommand	Flags	Description
mode get	—	Whether this mount's project is <code>open</code> (default-readable) or <code>protected</code> (default-denied, entries need an explicit grant).
mode set <open protected>	—	Switch project ACL mode. Admin-or-owner-gated: exit 2 on <code>forbidden / not_found</code> . Switching <code>to protected</code> additionally needs the org's plan to include folder ACLs — exit 2 with <code>plan_restricted</code> (402) if not.
audit [--entry-id] [--limit] [--offset]	—	The ACL change log — every grant/revoke, project- or entry-scoped.
entry get <entry-id>	—	The ACL roster on one entry (direct + inherited rows). <entry-id> accepts a mount path too (resolved first). Admin-gated: 403s for a member who is neither an org admin nor an entry admin (exit 2), since the roster reveals who else has access.
entry set <entry-id>	--principal <id>, --permission <read comment comment_upload write admin>, --effect <allow deny> (default allow), --inherit	Add or update a grant. Creating one needs the org's plan to include folder ACLs — exit 2 with <code>plan_restricted</code> (402) on Starter.
entry remove <entry-id> <principal-id>	—	Revoke a principal's grant on an entry.
principals [--query] [--limit] [--offset]	—	Look up an ACL <code>principal id</code> — the value <code>entry set --principal</code> and <code>mr assign</code> take. <code>--query</code> searches display name/email case-insensitively; without it, lists the org's principals. One page is capped at 50 rows; <code>--offset</code> pages past it, for a search as well as the full list.

```

swarmfile acl mode get
swarmfile acl mode set protected
swarmfile acl audit --entry-id 3f9a2b --limit 50
swarmfile acl principals --query "Priya"
swarmfile acl entry get 3f9a2b
swarmfile acl entry set 3f9a2b --principal u_9f2a --permission write --inherit
swarmfile acl entry remove 3f9a2b u_9f2a

```

`acl entry` is the per-entry grant surface the project-wide `mode` toggle doesn't cover; the same grants are editable from the dashboard's permissions panel, and (on Windows) projected as NTFS DACLS on the mount.

e2e-key

End-to-end encryption device keys.

Subcommand	Description
status	This device's own enrolled key — id and public key — read straight from the local cache file, no hub round trip. <code>enrolled: false</code> means this device hasn't finished enrolling yet (or the binary was built without the <code>e2e</code> feature at all, flagged with <code>e2eBuild: false</code>).
list	Every device key on this account — every machine, not just this one. For finding the id of a lost device to revoke.
revoke <id>	Revoke a device key by id (see <code>list</code>) — e.g. after a lost laptop.

There's no `enroll` here — a device enrolls its own key automatically at startup.

```

swarmfile e2e-key status
swarmfile e2e-key list
swarmfile e2e-key revoke key_9f2a

```

device-transfer

Move a device's end-to-end project keys onto a second device without going through the org's recovery-key quorum — a one-time, ten-minute code links the new device to an already-enrolled one and moves the wrapped project keys across, entirely client-side. The private key never touches the hub.

Subcommand	Description
start	Run on the device that already has the project's key. Prints a one-time code, waits for a new device to link, then uploads this mount's currently-mounted project key to it automatically. Requires an enrolled E2E key and a mounted E2E project on this engine.
link <code> [--label]	Run on the new device with the code from <code>start</code> . Waits for the other device to upload, then finalizes — re-wrapping every recovered project key to this device's own already-enrolled key. <code>--label</code> names the device (shown later in <code>e2e-key list</code> and Settings → My Devices).

Both commands poll every couple of seconds until the handshake completes or the code expires; past 10 minutes either side reports a clear "expired, run `start` again" rather than hanging.

```
# On the device that already has the project:
swarmfile device-transfer start
# Transfer code: 7F3K-9QRT
# On the new device, within 10 minutes, run:
#   swarmfile device-transfer link 7F3K-9QRT

# On the new device:
swarmfile device-transfer link 7F3K-9QRT --label "Jordan's laptop"
```

`start` transfers only the mount's currently-configured project — run it again, once per mounted project, on a device that holds more than one. There's no `--enroll` step on the new device: every engine enrolls its own long-term device key automatically at startup, and `link` / `finalize` re-wrap to that key directly.

If nobody's engine is online to complete an ordinary access grant, an org member can also fulfill it manually from the web dashboard rather than running this at all — see [End-to-End Encryption Setup](#).

admin nodes

The hub's P2P mesh node allow-list — owner-only. `node_id` is the 64-hex-char Ed25519 public key an engine self-generates (visible in the engine's own startup log, or in a not-yet-authorized node's connection-refused message). The same list, with an authorize form and a revoke button, is also available from **Settings → Nodes** in the web dashboard.

Subcommand	Flags	Description
<code>list</code>	—	Every node currently on the org's allow-list.
<code>authorize <node-id> <user-id> [--label]</code>	—	Add a node, authenticating as <code>user-id</code> .
<code>revoke <node-id></code>	—	Remove a node — it can no longer join the mesh.

```
swarmfile admin nodes list
swarmfile admin nodes authorize a1b2c3...64hex user_88 --label "render-farm-3"
```

admin retention

The organization's trash/history retention window and its **size-adaptive floor**, plus the history-retention default for NEW projects — owner-only. When a project approaches its storage limit, the hub automatically shrinks the retention window under pressure; the floor is the tightest it may ever shrink to. Both are `null` (the global defaults) until set. `keepFullHistoryDefault` decides what newly created projects start with: `true` (the product default) keeps every saved version, `false` starts them on named commits only — a creator's explicit choice at creation still wins. The same controls are in **Settings → Organization** in the web dashboard.

Subcommand	Flags	Description
<code>get</code>	—	The current window, floor, and new-project retention default.
<code>set</code>	<code>--days <1-3650></code> , <code>--floor <1-3650></code> , <code>--clear-days</code> , <code>--clear-floor</code> , <code>--keep-full-history-default <true false></code>	Set and/or clear the day values, and/or set the new-project default. Omitted fields are left unchanged. The hub rejects a floor above the window (it would be inert).

```
swarmfile admin retention get
swarmfile admin retention set --days 120 --floor 30
swarmfile admin retention set --clear-floor
swarmfile admin retention set --keep-full-history-default false
```

admin rate-limits

The organization's **request budgets**: the per-org and per-user request ceilings, in requests per minute — owner/admin-configurable. The same limits appear in the tray's **Settings → Developer** panel and the web dashboard's **Settings → Organization** page. `show` reports the current values alongside each plan default and maximum; `set` changes them. The Free plan's limits are fixed, Starter/Pro cap at their tier maxima, and Enterprise's org ceiling can be removed entirely. A PAT shares its owner's per-user window; an API key gets its own window at the plan's default per-user value; content streaming isn't counted here.

Subcommand	Flags	Description
<code>show</code>	—	The current per-org and per-user ceilings, their plan defaults and maxima, and whether each is configurable.
<code>usage</code>	—	This minute's usage against the ceilings, for the whole org and for you (the caller), including the separate media-read windows.
<code>set</code>	<code>--org <n></code> , <code>--user <n></code> , <code>--clear-org</code> , <code>--clear-user</code> , <code>--unlimited-org</code> , <code>--for <dur></code>	Set and/or clear either ceiling. Omitted fields are left unchanged. <code>--clear-*</code> resets a field to the plan default; <code>--unlimited-org</code> removes the org ceiling entirely (Enterprise only); <code>--for 6h</code> (units <code>s / m / h / d</code> , max 30 days) makes a raise auto-revert. The hub rejects a value above the plan's max and refuses any change on the Free plan.

`rate-limits` is also available at the top level (`swarmfile rate-limits ...`) — the same subcommands, without the `admin` prefix.

```

swarmfile admin rate-limits show
swarmfile admin rate-limits usage
swarmfile admin rate-limits set --org 4000 --user 450
swarmfile admin rate-limits set --org 8000 --for 6h
swarmfile rate-limits set --clear-org
swarmfile rate-limits set --unlimited-org

```

notifications

This account's notification inbox — comments, mentions, unlock requests, merge requests, RFIs, quarantine, failed uploads, and more. `list / unread-count / read / read-all` hit the **hub's full notification feed** (the same inbox the dashboard's bell shows), not just this machine's own upload failures.

Subcommand	Flags	Description
<code>list</code>	<code>--kind</code> , <code>--unread</code> , <code>--since <time></code> , <code>--limit</code>	Notification history, newest first. <code>--since</code> accepts Unix seconds, epoch milliseconds, a duration (<code>2h</code> , <code>7d</code>), or an RFC 3339 timestamp. <code>--kind</code> accepts a hub notification kind (the same open vocabulary — comments, mentions, unlocks, RFIs, MRs, quarantine, uploads, storage, ...); an unknown value is rejected with the accepted list and exit 2, so a typo can't silently return the whole inbox.
<code>unread-count</code>	—	Cheap count for a shell prompt or a CI/cron check.
<code>read <id></code>	—	Mark one notification read.
<code>read-all [--kind]</code>	—	Mark every notification read, optionally narrowed to one kind. An unknown <code>--kind</code> is rejected (exit 2) rather than marking the whole inbox read.
<code>preferences get</code>	—	Current email/webhook delivery preferences — the webhook status (<code>url / enabled / configured</code>) rides along in this same output.
<code>preferences set</code>	<code>--email-enabled [<true false>]</code> (bare = on), <code>--quiet-start / --quiet-end (HH:MM)</code> , <code>--timezone</code> , <code>--kind <KIND=on off default></code> (repeatable)	Update one or more fields — omitted flags leave the existing value untouched. <code>--quiet-start / --quiet-end</code> must both be set together or both left unset. Clearing quiet hours or the timezone is dashboard-only: the API takes an explicit null for it, but the CLI has no null form (an empty value is a validation refusal), so there is no CLI clear. <code>--kind</code> sets a per-kind email override (<code>off</code> mutes that event, <code>on</code> forces it, <code>default</code> clears that kind's whole override — email and webhook — back to the kind's defaults: the global switches plus that kind's own default, where one applies); the CLI merges into your current overrides, so other kinds and their webhook flags are preserved. An unknown kind or malformed spec exits 2. Exit 2 on a <code>400</code> (mismatched pair or malformed time).
<code>preferences webhook set <url></code>	<code>--enabled [<true false>]</code> (bare = on)	Set (or change) YOUR OWN personal webhook — a third notification channel next to the in-app bell and email, one URL per user. The <code><url></code> argument is required on every call (no partial "just flip <code>--enabled</code> " shorthand — pass the same URL again to toggle it alone). Mints a fresh signing secret whenever the URL changes, or is set for the first time, and prints it ONCE. Distinct from the org-wide <code>swarmfile webhook</code> command group below: this one is fire-and-forget, with no delivery log or auto-disable.
<code>preferences webhook clear</code>	—	Remove your personal webhook.
<code>preferences webhook rotate-secret</code>	—	Mint a fresh signing secret without changing the URL. Hard cutover — the old secret stops verifying immediately. Prints the new secret ONCE.
<code>preferences webhook test</code>	—	Fire one synthetic event at your configured webhook right now, synchronously. Nothing is logged — this channel has no delivery log — the result prints directly.

```

swarmfile notifications list --unread --limit 20
swarmfile notifications list --kind comment_mention
swarmfile notifications read-all
swarmfile notifications preferences set --quiet-start 22:00 --quiet-end 07:00 --timezone America/New_York
swarmfile notifications preferences set --kind file_changed_by_other=off
swarmfile notifications preferences set --kind comment_added=default # clears email + webhook for that kind

```

share

Share links — a read/comment/respond link to one file, handed to someone outside the org. This covers only the member-authed management side: creating your own links, listing/revoking them, and (owner-only) the access log and email block-list. The anonymous recipient's own side of a link is the web-hosted `/s/:token` page, not a CLI surface.

Subcommand	Flags	Description
create <entry-id>	<code>--expires-at <unix-secs></code> , <code>--max-accesses <N></code> , <code>--password <string></code> , <code>--allow-comments</code> , <code>--allow-responses</code>	Create a link for one file. <code><entry-id></code> accepts a mount path too (the path is resolved first). Exactly one of <code>--allow-comments</code> / <code>--allow-responses</code> is meaningful — the hub picks the more permissive one if both are set. <code>--allow-responses</code> is the RFI share-link path (needs a directory entry, not a file) — see RFI attachments . Exit 2 on a <code>400</code> (missing/invalid <code>entry-id</code> , or wrong entry type for the requested permission) or a <code>404</code> (entry gone).
list	<code>--entry-id</code>	The caller's own share links, optionally narrowed to one file.
revoke <id>	—	Revoke a link. Exit 2 on <code>404</code> (not found, or not yours).
access-log <id>	<code>--limit</code> , <code>--offset</code>	Who has actually opened this link (the mandatory-email-gating log, not just a hit count). Owner (creator) only — exit 2 on <code>404</code> .
block <id> <email>	—	Block one email from this link going forward. Owner only. Takes effect on that email's next gated request — an already-issued 30-day session cookie isn't separately revoked.
unblock <id> <email>	—	Unblock a blocked email address. Owner only.

```
swarmfile share create e-1a2b3c --allow-comments --expires-at 1735689600
swarmfile share list --entry-id e-1a2b3c
swarmfile share access-log shr_9f2a --limit 20
swarmfile share block shr_9f2a spammer@example.com
swarmfile share revoke shr_9f2a
```

activity

The org/project activity feed — commits, branch switches, comments, CI runs, membership/ACL changes, who did what and when. An org owner sees everything; everyone else only what they have read access to (owner-only sources like quarantine and org-policy events project to empty for them).

Subcommand	Flags	Description
list	<code>--since / --until <time></code> , <code>--limit</code> , <code>--kind <csv></code> , <code>--project-id</code> , <code>--entry-id</code> , <code>--actor-id</code> , <code>--cursor</code>	The feed, newest-filtered-window first. Defaults to the last 7 days when <code>--since / --until</code> are omitted. <code>--since / --until</code> accept Unix seconds, epoch milliseconds, a duration (<code>2h</code> , <code>7d</code>), or an RFC 3339 timestamp. <code>--kind</code> takes a comma-separated list of event kinds. <code>--cursor</code> resumes from a previous response's pagination cursor.
export	<code>--since / --until</code> , <code>--kind</code> , <code>--project-id</code> , <code>--out <path></code>	Unbounded CSV export (capped at 100k rows hub-side) — org-owner-only , and needs the Pro+ <code>audit_log</code> entitlement (exit 2 on the <code>403</code> either gate returns). Prints CSV to stdout by default; <code>--out</code> writes it to a file instead. Not paginated — narrow <code>--since / --until</code> if you need a complete window (a response is capped at 100k rows).
tail	<code>--interval <secs></code> (default 3), <code>--since <time></code> , <code>--kind</code> , <code>--project-id</code> , <code>--actor-id</code>	Follow the feed: print new events as they land, until Ctrl+C. Starts from now unless <code>--since</code> is given (same time shapes as <code>list</code>).

`tail` follows by polling the same `GET /activity` endpoint `list` uses — polling is the supported way to follow the feed. Expect a few seconds of latency, bounded by `--interval`. With `--json` it prints one event per line (JSON-lines), so a consumer can process the stream incrementally instead of re-parsing a document that repeats each poll.

```
swarmfile activity list --kind mr_merged,commit_created --limit 50
swarmfile activity list --since 1735689600 --project-id p_9f2a
swarmfile activity tail --kind commit_created
swarmfile activity export --since 1767225600 --until 1798675200 --out audit-2026.csv
```

webhook

Manages this **org's** [outbound webhook](#) subscriptions — signed HTTP callbacks fired on file/branch/merge/ACL/membership events, straight into Slack, Zapier, a SIEM, or your own endpoint. Everything the web dashboard's webhook settings page does is available here too; this is the org-wide subscription list, distinct from the single personal channel under `notifications preferences webhook` above.

Subcommand	Flags	Description
<code>list</code>	—	Every webhook subscription for this mount's org. Never shows signing secrets.
<code>create <url></code>	<code>--event <kind></code> (repeatable), <code>--description</code>	Subscribe a URL to org activity events. Prints the raw signing secret ONCE — copy it now, the hub never returns it again. <code>--event</code> may be repeated to subscribe to specific kinds only; omit it to receive every kind.
<code>delete <id></code>	—	Remove a subscription. Immediate.
<code>reenable <id></code>	—	Clear a subscription's auto-disable — the hub disables one automatically after too many consecutive delivery failures.
<code>rotate-secret <id></code>	—	Mint a fresh signing secret. Hard cutover — the old secret stops verifying immediately. Prints the new secret ONCE.
<code>update <id></code>	<code>--url</code> , <code>--description</code> , <code>--clear-description</code>	Update a subscription's URL and/or description. An omitted flag keeps its current value. Doesn't touch which event kinds it receives or its custom headers.
<code>deliveries <id></code>	<code>--limit</code> , <code>--cursor</code>	A subscription's delivery log, newest first, cursor-paginated (use the previous response's <code>nextCursor</code>).
<code>test <id></code>	—	Fire one synthetic event at a subscription's URL right now, synchronously. Recorded in the delivery log, but never counts toward auto-disable health tracking.
<code>retry-delivery <id></code> <code><delivery-id></code>	—	Manually re-send one past delivery. A delivery with no stored payload can't be retried.

```

swarmfile webhook create https://hooks.example.com/swarmfile --event mr_merged --event commit_created
swarmfile webhook list
swarmfile webhook deliveries wh_9f2a --limit 20
swarmfile webhook rotate-secret wh_9f2a

```

Scripting against the CLI

Pass the global `--json` (before or after the subcommand — both work) and it prints the engine's raw response, which is the intended path for automation — commit from CI, hydrate before a render, check `conflicts` in a pre-flight script (`conflicts resolve <path> --auto` clears text conflicts headlessly, exiting 2 when a human is needed). A usage error under `--json` emits `{"code":"usage_error",...}` rather than clap's prose, so the same parser handles both. Pair it with an API key and no part of the loop needs a human.

For reacting to changes rather than polling for them, `tail` streams this project's activity as it happens — hold that connection open for sub-second reaction. If your integration doesn't want to hold a connection open at all, `swarmfile webhook create` (org owner) or `notifications preferences webhook set` (personal) both set up a signed HTTP callback to a URL you control instead. They differ in delivery: org webhooks are delivered on a periodic schedule, not in real time, with retries and a delivery log; the personal channel is fire-and-forget, enqueued as each notification happens, with no delivery log, retry, or auto-disable.

swarmfile-doctor

`swarmfile-doctor` diagnoses Swarmfile connectivity: DNS, hub, cloud storage, the peer-to-peer network, mDNS, and OIDC — plus `install-layout` (detects an install an update can't reach) and a `plan` check. The `plan` check runs on every machine: it reads your organization's plan standing, and it **fails when Swarmfile isn't signed in**, because a machine that can't read your plan generally can't reach your files either. On seed/NAS nodes it additionally verifies your plan includes the seeding entitlement. Signed in as a guest reviewer, it passes and says so — guests aren't members, so there's no plan to read. It's a standalone binary — unlike `swarmfile doctor`, which asks a running engine to run its probe suite, `swarmfile-doctor` needs no running engine at all. Use it on a machine where the engine isn't up, won't start, or you're not sure which is true. It ships on `PATH` with every desktop install, and — uniquely among these tools — is also published as a standalone download for machines where installing the app isn't an option: Developer ID-signed and notarized on macOS, unsigned on Windows and Linux (so the Windows SmartScreen prompt is expected — see [Getting Started](#)); see also [Security](#).

Flags

Flag	Description
<code>--json</code>	Machine-readable output
<code>-q, --quiet</code>	Suppress non-error output. Implies <code>---json</code> for stdout.
<code>--report <path></code>	Custom path for the JSON report. A report is always written, <code>--json / --repair</code> or not — this only overrides where. Default <code>~/.cache/swarmfile/last-doctor.json</code> (macOS/Linux) or <code>%USERPROFILE%\Swarmfile\last-doctor.json</code> (Windows).
<code>--hub-url <url></code>	Override the hub URL to check. Without it, the install's configured hub is used: <code>\$SWARMFILE_HUB_URL</code> , else the hub the install's <code>config.json</code> names, else the build's fallback (the hosted hub for a release build).
<code>--no-iroh</code>	Skip the peer-to-peer connectivity probe
<code>--no-mdns</code>	Skip the mDNS discovery probe
<code>--iroh-peer <HOST:PORT></code>	Peer-to-peer network peer to probe against (default <code>192.0.2.1:4433</code>)
<code>--cache-dir <path></code>	Engine cache dir. Default <code>~/.cache/swarmfile</code> (macOS/Linux), <code>%LOCALAPPDATA%\swarmfile</code> (Windows).
<code>-v, --verbose</code>	Verbose output
<code>--repair</code>	Download the full installer for this platform and run it, replacing every component.
<code>--yes</code>	Skip the confirmation prompt for <code>--repair</code>

Exit codes

The process exit code is the machine-readable verdict, so a script or provisioning step can gate on it without parsing the report:

Code	Meaning
0	Every check passed.
1	At least one check failed .
2	At least one warn , and no fails.

`--repair`

`--repair` is the most consequential flag here: it downloads the full installer for the current platform and runs it, replacing every Swarmfile component on the machine. It's the headless equivalent of the Desktop App's Diagnostics -> Reinstall button, and it's meant for machines that are broken badly enough that a targeted fix isn't practical — not for routine troubleshooting.

By default `--repair` stops for a confirmation prompt before it does anything. Pass `--yes` alongside it to skip that prompt for unattended or scripted repair (for example, a provisioning step that always wants a clean install). See [Operations](#) for where repair fits into a broader recovery workflow.

Examples

Plain diagnostic run:

```
swarmfile-doctor
```

JSON output suitable for attaching to a support ticket:

```
swarmfile-doctor --json --report doctor-report.json
```

Unattended repair, e.g. from a provisioning script:

```
swarmfile-doctor --repair --yes
```

swarmfile-migrate

`swarmfile-migrate` bulk-imports data you already have into a Swarmfile project: a local directory, a flat list of file paths, or an S3-compatible bucket. It's a standalone binary — no running engine required — meant for the first-time import of an existing archive, not for day-to-day file operations. Every desktop installer carries it, and the Linux `.deb`, macOS `.pkg` and Windows installers put it on `PATH` (`/usr/bin` on Linux, `/usr/local/bin` on macOS, or the Windows install folder); a macOS `.dmg` -only install needs one `.pkg` (or **Diagnostics → Reinstall**) run for that entry — so if the desktop app is installed there's nothing extra to download.

This page is the flag reference. For a practical walkthrough with a worked example, see [Migrating Existing Data](#).

Run with `--dry-run` first. It walks the source and reports what would be migrated without uploading anything, which is the safe way to check your exclude patterns and source/destination settings before committing to a real run.

Source selection

Pick exactly one source: a local directory, a flat file list, or an S3-compatible bucket.

Flag	Description
<code>--source</code>	Local directory source
<code>--source-list</code>	Flat file-list source
<code>--source-root</code>	Root to resolve relative paths in <code>--source-list</code> against
<code>--source-s3-bucket</code>	S3-compatible source bucket
<code>--source-s3-prefix</code>	Prefix within the S3 bucket to migrate
<code>--source-s3-endpoint</code>	S3-compatible endpoint URL
<code>--source-s3-region</code>	Default <code>us-east-1</code>

S3 support ships in every released installer — `--source-s3-*` works out of the box, with no extra download. The AWS SDK it needs is large and lives in the `swarmfile-migrate` tool rather than the Swarmfile engine; a from-source build can opt out with `--no-default-features`, and a binary built that way fails immediately (naming the missing feature) instead of partway through a migration.

Destination

Flag	Description
<code>--root-path</code>	Destination path inside the project. Default <code>/</code>
<code>--hub-url</code>	Your hub URL, e.g. <code>https://hub.swarmfile.com</code> (or your self-hosted hub's URL). Defaults to <code>\$SWARMFILE_HUB_URL</code> if set, otherwise the hosted hub (<code>https://hub.swarmfile.com</code>) — so a machine whose environment already names its hub does not need this flag
<code>--org-id</code>	Target organization
<code>--project-id</code>	Target project. Defaults to <code>\$SWARMFILE_PROJECT_ID</code> if set; with neither, entries are created project-less — invisible to any mount that scopes its feed to a project (which every engine does by default), so set one for a normal import
<code>--oidc-issuer</code>	OIDC issuer for authenticated hub access
<code>--oidc-client-id</code>	OIDC client ID
<code>--oidc-refresh-token</code>	OIDC refresh token

Authentication

The simplest and primary way to authenticate is a **project-scoped API key** in the environment: set `SWARMFILE_API_KEY=sf_key_...` and `swarmfile-migrate` picks it up automatically — the same headless credential every other Swarmfile binary honors — with no auth flags at all. This is the recommended path for a CI runner or a render-farm import box.

The `--oidc-issuer` / `--oidc-client-id` / `--oidc-refresh-token` flags are the alternative: pass **all three together** to authenticate with an OIDC refresh token instead. If you pass no API key and not all three OIDC flags, hub calls go out unauthenticated and will fail against any hub that enforces auth.

Exclusions

Flag	Description
<code>--exclude</code>	Glob pattern to exclude. Repeatable. Matching a directory prunes its whole subtree.
<code>--exclude-from</code>	File containing exclude patterns, one per line

Reliability

Flag	Description
<code>--concurrency</code>	File-level concurrency. Default <code>8</code>
<code>--block-concurrency</code>	Max concurrent block uploads, shared across all files. Default <code>4</code>
<code>--retry</code>	Max retry attempts per file — both in-run (a transient failure retries immediately, with backoff, before the run moves on to the next file) and across separate re-inocations of this tool against the same <code>--state-db</code> . Default <code>3</code> ; <code>0</code> disables retries either way
<code>--verify</code>	Verify each migrated file against the source with a CID/hash comparison. On by default (<code>true</code>); pass <code>--verify=false</code> to skip — verifying means reading the source a second time, so skipping it roughly halves I/O for a very large migration where you trust the transport
<code>--state-db</code>	Local state DB path. Tracks what's already uploaded and enables resuming a re-run. Default <code><source>/swarmfile-migration/migrate.db</code> ; for a <code>--source-list</code> source, it sits alongside the list file instead
<code>--report-json</code>	Write a JSON report of the run
<code>--start-at</code>	Resume marker to start from

Safety

Flag	Description
<code>--dry-run</code>	Report what would be migrated without uploading anything
<code>--force</code>	Re-migrate files already registered/uploaded in a prior run, instead of skipping them

Encryption & erasure coding

There's no `--encrypt` flag — migrated data is encrypted the same way any other write to the project is, driven by environment variables rather than a migration-specific setting: `SWARMFILE_ENCRYPTION_KEY` / `SWARMFILE_ENCRYPTION_PROJECT_ID`, or simply `--project-id` combined with OIDC credentials, resolves the key. Encryption is on by default; if no key source is configured for the invocation, the tool warns and migrates the data as plaintext rather than refusing to run — unless encryption was explicitly requested, in which case it fails instead. Erasure coding is controlled by `SWARMFILE_EC_ENABLED` (`on` / `off` / `auto`): default `auto` (also what an unset or unrecognized value falls back to) **never erasure-codes during migration**, because `auto` is evaluated from live peer/network topology and a bulk import has none — set `SWARMFILE_EC_ENABLED=on` (or `true`) to force erasure coding on the migrated blocks.

Examples

Step 1 — preview with `--dry-run`. Always start here: it walks the source and reports what would be migrated without uploading anything, so you can confirm your exclude patterns and source/destination settings before committing to a real run:

```
swarmfile-migrate --source /Volumes/nas/ProjectArchive \
  --org-id acme-films --project-id feature-01 \
  --exclude "*.cache" --exclude "**/tmp/**" \
  --dry-run
```

Step 2 — the real run. Once the dry-run looks right, drop `--dry-run` (post-upload verification is on by default, so no flag needed). Here `SWARMFILE_API_KEY` supplies the credential and `$SWARMFILE_HUB_URL` the hub, so neither needs a flag:

```
export SWARMFILE_API_KEY=sf_key_...
swarmfile-migrate --source /Volumes/nas/ProjectArchive \
  --org-id acme-films --project-id feature-01 \
  --exclude "*.cache" --exclude "**/tmp/**" \
  --state-db ./migrate-state.db
```

S3-compatible bucket (preview it with `--dry-run` the same way first):

```
swarmfile-migrate --source-s3-bucket project-archive \
  --source-s3-prefix renders/2025 \
  --source-s3-endpoint https://s3.us-west-000.backblazeb2.com \
  --hub-url https://hub.swarmfile.com \
  --org-id acme-films --project-id feature-01 \
  --dry-run
```

swarmfile-search

`swarmfile-search` runs full-text search across the local Swarmfile metadata cache, with fallback to the hub. It's a standalone binary for searching from a terminal or a script, rather than the dashboard omnibox. Every desktop installer carries it, and the Linux `.deb`, macOS `.pkg` and Windows installers put it on `PATH` (`/usr/bin` on Linux, `/usr/local/bin` on macOS, or the Windows install folder); a macOS `.dmg` -only install needs one `.pkg` (or **Diagnostics → Reinstall**) run for that entry.

The local cache is entries-only: it indexes file and folder names in a local full-text index, and it's what a plain query checks first. Project and people (principal) name matches only come from the hub fallback — the local cache doesn't mirror those, so a query for a project or person name needs a hub round-trip (or `--remote`) to find anything. On the hub the two halves also live in different places: entries are indexed per project, while project and people names are org-wide. A query scoped with `--project` (or `$SWARMFILE_PROJECT_ID`) therefore asks both routes and merges them, so you still get project/people matches alongside the file hits; an unscoped hub query searches project and people names only — **file hits from the hub need `--project`**, or a warm local cache. Either way, it's a metadata search — it finds things by name, not by file contents.

For the conceptual overview of how search resolves results and where else you can search from, see [Search](#).

Usage

```
swarmfile-search <query>
```

Flags

Flag	Description
<code>--project</code>	Scope the search to a single project by ID. Defaults to <code>\$SWARMFILE_PROJECT_ID</code> when set
<code>--limit</code>	Maximum results. Default <code>20</code>
<code>--remote</code>	Force the search against the hub. Conflicts with <code>--local-only</code>
<code>--local-only</code>	Force the search against the local cache only. Conflicts with <code>--remote</code>
<code>--json</code>	Print raw JSON instead of a human-readable list. Emits the hub's <code>{files, folders, projects, principals}</code> shape plus a top-level <code>source</code> field

Examples

Plain query across the local cache, falling back to the hub if needed:

```
swarmfile-search "concept_v3"
```

Scoped to a project, capped at 5 results, local cache only:

```
swarmfile-search "elevation" --project proj_xyz789 --local-only --limit 5
```

JSON output for scripting:

```
swarmfile-search "plan.dwg" --json
```

Alongside the hub-shaped `{files, folders, projects, principals}` object, `--json` adds a top-level `source` field so a script can tell where the results came from without changing the other keys — `local`, `local (no fallback)`, `remote`, or `remote (local fallback)`.

swarmfile-seed

`swarmfile-seed` seeds a single local file into the Swarmfile hub: it reads the file, chunks it, uploads the resulting blocks, and registers a file entry (creating any missing parent directories along the way) at a virtual path you specify. It's a standalone binary with no subcommands — one invocation seeds one file. Every desktop installer carries it, and the Linux `.deb`, macOS `.pkg` and Windows installers put it on `PATH` (`/usr/bin` on Linux, `/usr/local/bin` on macOS, or the Windows install folder); a macOS `.dmg` -only install needs one `.pkg` (or **Diagnostics → Reinstall**) run for that entry.

`swarmfile-seed` is a one-shot ingestion tool, not a persistent process. It uploads the file you point it at and exits. For running a machine as an ongoing, persistent participant in the peer-to-peer swarm, see [swarmfile seed enable](#) on the full client, covered in [Self-Hosted Seed Nodes](#).

Flags

Flag	Description
<code>--file</code>	Path to the local file to upload
<code>--path</code>	Virtual path in the filesystem (e.g. <code>/Projects/plan.dwg</code>)
<code>--hub-url</code>	Hub URL. Defaults to <code>\$SWARMFILE_HUB_URL</code> , else the hub the install's <code>config.json</code> names, else the build's fallback (<code>https://hub.swarmfile.com</code> in a release build, <code>http://localhost:8787</code> in a debug build). Override for a self-hosted / Enterprise hub
<code>--cdc</code>	Use content-defined chunking (FastCDC) instead of fixed-size
<code>--auto-profile</code>	Auto-select chunking strategy based on file extension (AEC profiles). Conflicts with <code>--cdc</code>

`--file` and `--path` are both required. `--path` must not be empty or contain empty segments (no leading/trailing/double slashes past trimming) — it's validated before any I/O happens.

Chunking

By default, `swarmfile-seed` chunks the file with a fixed chunk size. `--cdc` switches to content-defined chunking (FastCDC), which tends to deduplicate better across files with small internal shifts. `--auto-profile` instead picks a chunking strategy per the file's extension using the same AEC (architecture/engineering/construction) profiles the rest of the ingest pipeline uses, and is mutually exclusive with `--cdc`. If neither flag is passed, fixed-size chunking is used.

What it does not do

Unlike `swarmfile-migrate`, `swarmfile-seed` uploads exactly one file per invocation — there's no directory-walk or file-list mode. It also does not erasure-code the upload — a deliberate scope limit for this single-file admin/seeding tool.

There is no `--project-id` or `--org-id` flag; project selection is environment-driven. The entry has to land in a project — the project-scoped metadata route is the only one that can accept a create — so `swarmfile-seed` resolves it from `SWARMFILE_PROJECT_ID` (falling back to `SWARMFILE_ENCRYPTION_PROJECT_ID` when that is the only one set, for the managed-key path) and fails with an actionable error before uploading anything if neither is set. If both are set they must name the same project. Any encryption-key lookup keyed to a project is likewise driven entirely by environment variables rather than a CLI flag. The key sources are:

- `SWARMFILE_ENCRYPTION_KEY` — a static hex-encoded key you supply directly.
- `SWARMFILE_ENCRYPTION_PROJECT_ID` — resolve a managed-tier project's hub-held key. This path additionally needs OIDC credentials in the environment: `SWARMFILE_OIDC_ISSUER`, `SWARMFILE_OIDC_CLIENT_ID`, and `SWARMFILE_OIDC_REFRESH_TOKEN` (all three).

Encryption is **on by default**. If no key source is configured for the invocation, the tool warns and seeds the file as plaintext rather than refusing to run — unless encryption was explicitly requested (any of the above set), in which case it fails instead. To silence the plaintext-mirror warning on a node that is deliberately seeding without encryption, set `SWARMFILE_ENCRYPTION_ENABLED=false`.

Example

```
SWARMFILE_PROJECT_ID=<project-id> swarmfile-seed --file ./plan.dwg \
--path "/Projects/Site A/plan.dwg" \
--hub-url https://hub.swarmfile.com --auto-profile
```

swarmfile-brlock

`swarmfile-brlock` drives the engine's byte-range lock API directly. It talks to a **running engine** over its control socket, so start the engine (or the Desktop App) first. It's a developer/integrator tool — a test client for the same lock primitive that a native worksharing app plugin (the kind of Revit-class CAD/BIM integration that needs to lock only the elements one person is editing, not an entire file) would implement against. It is not something a typical end user runs day to day, and **no installer ships it**: build it from a checkout when you're working on an integration (`cargo build -p swarmfile-engine --bin swarmfile-brlock`).

For the conceptual introduction to byte-range locking and how it differs from a whole-file entry lock, see [Working with Files](#).

Global flags

Flag	Description
<code>--cache-dir</code>	Engine cache dir to connect against. Defaults to <code>\$SWARFFILE_CACHE_DIR</code> , else <code>~/.cache/swarmfile</code> (macOS/Linux) or <code>%LOCALAPPDATA%\swarmfile</code> (Windows).

acquire

Acquire an exclusive byte-range lock.

Flag	Description
<code>--entry-id</code>	Entry to lock a byte range on
<code>--offset</code>	Start offset, in bytes (<code>u64</code>). Default <code>0</code>
<code>--length</code>	Length of the range, in bytes (<code>u64</code>). Default <code>0</code> , which means "to end of file," not a zero-length range
<code>--handle-ref</code>	Handle the lock is tied to

release-handle

Release every lock tied to a handle ref.

Flag	Description
<code>--handle-ref</code>	Handle whose locks should all be released

Example: acquire and release a range

A plugin integration would typically acquire a range around the elements being edited, hold it for the duration of the edit, and release everything tied to its handle on close — this simulates that sequence from the command line:

```
swarmfile-brlock acquire --entry-id ent_9f2a --offset 4096 --length 512 --handle-ref plugin-session-1

# ...edit is in progress, range is locked...

swarmfile-brlock release-handle --handle-ref plugin-session-1
```

Releasing by handle ref rather than by individual range means a plugin doesn't need to track every range it acquired — closing out a session releases everything that session took.

Exit codes

Meaningful for scripting a CI/integration test against a real lock denial, which is the tool's main use case:

Exit code	Meaning
<code>0</code>	Success.
<code>2</code>	The lock was denied (a conflicting range is already held).
<code>1</code>	Any other error.

swarmfile-runner

`swarmfile-runner` runs the engine in **runner mode**: it sets `SWARMFILE_RUNNER_MODE=true` (when the variable isn't already set), then re-execs the sibling `swarmfile-engine` binary with the same arguments and standard I/O. It has no subcommands or flags of its own — every engine flag passes straight through — and it runs the real engine boot path (config, single-instance lock, lifecycle) exactly as `SWARMFILE_RUNNER_MODE=true swarmfile-engine` would.

It exists so the headless CI role reads as its own tool rather than "the engine with an env var". For what a runner does and how to wire it into a CI system, see [Runner \(Headless CI\)](#).

Because it re-execs the engine, `swarmfile-runner --help` is the engine's help. The pass-through flags:

Flag	What it does
<code>--build-info</code>	Print the enabled cargo features and target architecture, then exit.
<code>--status</code>	Engine, hub, peers, and pending work at a glance (works while another engine holds the lock).
<code>--capabilities</code>	The compiled feature set of this binary.
<code>--print-env-vars</code>	The resolved <code>SWARMFILE_*</code> values this build consumed, secrets redacted.
<code>--pause</code> / <code>--resume</code>	Pause or resume sync on a running engine without unmounting.
<code>--quit</code>	Ask the running engine to drain and stop.
<code>--pin</code> / <code>--unpin</code> / <code>--list-pins</code> / <code>--verify-pins</code>	Offline cache pins: add, release, list, and verify.
<code>--prepare-offline [SECS]</code> / <code>--return-from-offline</code>	The engine-level Pack & Go pair — see Working Offline .

See [Engine Config File](#) → [Engine command-line flags](#) for the full descriptions.

Environment

Variable	Description
<code>SWARMFILE_RUNNER_MODE</code>	Forced to <code>true</code> when unset. An explicit value is preserved, so <code>SWARMFILE_RUNNER_MODE=false swarmfile-runner</code> still runs in normal mode
<code>SWARMFILE_API_KEY</code>	Project API key the runner authenticates as
<code>SWARMFILE_PROJECT_ID</code>	Project the runner serves
<code>SWARMFILE_BRANCH</code>	Branch the runner tracks (defaults to the project's default branch)

Runner mode deliberately serves no control socket. A driverless CI host needs no FUSE — `materialize` and runner mode never call the mount path.

Acquiring it

- **Linux** — the `.deb` installs it at `/usr/bin/swarmfile-runner`, next to `swarmfile-engine` (the Docker output carries it too; ask us about an `.rpm`).
- **macOS** — it ships inside `Swarmfile.app/Contents/MacOS/swarmfile-runner`, next to the bundled engine, with a `/usr/local/bin/swarmfile-runner` symlink on a prod install.
- **Windows** — the MSI installs it next to `swarmfile-engine.exe` in the install folder, and the app-only zip carries it.
- **Anywhere else** — build it from source: `cargo build --release -p swarmfile-engine --bin swarmfile-runner`. It must sit in the same directory as `swarmfile-engine`, because it re-execs the sibling resolved next to its own executable.
- **Equivalent without the alias** — on any engine binary, `SWARMFILE_RUNNER_MODE=true swarmfile-engine` behaves identically.

Example

```
SWARMFILE_API_KEY=sf_key... SWARMFILE_PROJECT_ID=proj_xyz789 swarmfile-runner
```

Because it re-execs the engine, the process's exit code and output are the engine's. If the sibling `swarmfile-engine` is missing, the alias prints an error and exits 127.

swarmfile-verify-history

`swarmfile-verify-history` is a standalone verifier for a project's signed commit history. It walks the commit DAG from a branch's current `HEAD` back toward genesis, refetching each commit object from the hub, and recomputes every hash from the commit's own facts to prove the chain hasn't been tampered with: each parent link resolves and each stored hash matches what its facts actually hash to. Every desktop installer carries it, and the Linux `.deb`, macOS `.pkg` and Windows installers put it on `PATH` (`/usr/bin` on Linux, `/usr/local/bin` on macOS, or the Windows install folder); a macOS `.dmg` -only install needs one `.pkg` (or **Diagnostics → Reinstall**) run for that entry.

It also verifies the hub's **checkpoint signatures**. Every checkpoint the hub has published for the branch is fetched and its signature checked against a public-key ring — see the trust caveat below.

It does its own minimal HTTP GETs, with no dependency on a running engine or mount, so it can run on a CI box or an audit host.

Usage

```
swarmfile-verify-history \
  --hub-url https://hub.swarmfile.com/orgs/org.../projects/proj... \
  --org-id org... \
  --project-id proj... \
  --branch main \
  --header 'Authorization: Bearer <token>'
```

`--hub-url` is the **project-scoped** base: the verifier appends `/history/...` and `/metadata/...` paths to it, while `--org-id` / `--project-id` identify the org and project to verify.

Add `--json` for a machine-readable report.

Flags

Flag	Description
<code>--hub-url <URL></code>	Project-scoped hub base URL (required) — the project this run verifies, e.g. <code>https://hub.swarmfile.com/orgs/org.../projects/proj...</code> .
<code>--project-id <ID></code>	Project to verify (required).
<code>--org-id <ID></code>	Org id (required).
<code>--branch <NAME></code>	Branch to verify. Default <code>main</code> .
<code>--header "Key: Value"</code>	Extra raw header, repeatable — pass whatever auth header this deployment requires (for example, <code>Authorization: Bearer ...</code>). No auth scheme is hardcoded.
<code>--checkpoint-hash <HASH></code>	Stop walking when this commit is reached (inclusive), instead of continuing to genesis. Use when older history has been pruned server-side.
<code>--pin-key <PUBKEY FILE></code>	Pin the hub's checkpoint-signing Ed25519 public key out-of-band (64 hex, base64, or a file path). Repeatable for key rotation.
<code>--json</code>	Emit JSON instead of the human-readable report.

What it proves — and what it doesn't

- **Structural integrity** is real: every hash recomputes and every parent link resolves across the commits the hub served. That shows the chain is internally consistent; independent assurance that the hub is attesting to the same chain it originally recorded comes from the signed checkpoints below.
- **Signatures without `--pin-key`** are checked against the key ring fetched from the *same hub being audited* (trust-on-first-use). That proves the checkpoints were signed by whatever key that hub currently publishes — not that the key is the hub's real one. The report says so on every run.
- **With `--pin-key`** signatures are verified only against keys you supply out-of-band; the run fails if none is in the hub's ring, and a checkpoint signed by any other key fails. This is the mode that gives independent assurance.
- Where checkpoint signing isn't enabled, the key route answers `404`: unpinned that reads as "verifiable history isn't enabled here" (informational, exit OK); with `--pin-key` it is a failure.

Exit code is non-zero on any verification failure, so it is safe to gate CI on.

swarmfile-lfs-transfer

`swarmfile-lfs-transfer` is Swarmfile's **git-LFS custom transfer agent**. `git-lfs` invokes it — it is not a command you run by hand. Point `lfs.url` at your Swarmfile project and set `git-lfs` to use this agent as the transfer, and pushes ride it chunked and resumable (up to the 256 GiB object cap), while pulls use the plain resumable GET path.

It is the engine's own agent exposed as a standalone binary. The distributed entry point is the engine subcommand `swarmfile-engine lfs-transfer` (the engine binary already ships on every platform); the dedicated `swarmfile-lfs-transfer` binary is a dev/testing build that no installer ships — build it from a checkout with `cargo build -p swarmfile-engine --bin swarmfile-lfs-transfer`. Both share the same logic.

Don't confuse it with `swarmfile-lfs` — the standalone chunked-LFS agent that stock `git-lfs` points `lfs.url` at directly. `swarmfile-lfs-transfer` is the transfer *agent* `git-lfs` spawns; `swarmfile-lfs` is the higher-level agent and `reflect` tool, and it ships on `PATH` with every desktop install (alongside the separate Homebrew/cargo/ `curl` | `sh` channel for machines without the app).

How git-lfs calls it

`git-lfs` spawns a configured custom transfer agent with an operation and remote, and talks to it over stdio (the LFS transfer protocol). It takes no flags of its own. Configure it per project — point `git-lfs` at the shipped engine binary (or the dev-built dedicated binary, if you made one):

```
git config lfs.url "https://hub.swarmfile.com/orgs/<org>/lfs/<project>"
git config lfs.standalone-transfer-agent swarmfile-chunked
git config lfs.customtransfer.swarmfile-chunked.path swarmfile-engine
git config lfs.customtransfer.swarmfile-chunked.args lfs-transfer
```

See [git-LFS](#) for the full setup, the `.lfsconfig` auto-PR flow, and how LFS objects are billed and reflected onto the mount.

Notes

- LFS bytes are metered and **storage-billed like any other content** — there is no separate LFS SKU.
- Managed-tier LFS objects written through the block paths are encrypted like everything else; the direct presigned-PUT path stores plaintext even on managed (see the [git-LFS guide](#)), and E2E projects **refuse** LFS outright.
- Because it is spawned by `git-lfs`, there is nothing to run interactively; if you need a human-facing status, use the mount or `swarmfile status`.

git-remote-swarmfile

`git-remote-swarmfile` is the small **git remote helper** behind `git clone`, `git fetch` and `git push` on a `swarmfile://` URL. Git invokes it by name whenever it sees a `swarmfile://` remote, so you never run it yourself; it re-execs the engine binary in `git-remote` mode and needs no running engine, no drive and no mount.

Every desktop installer puts it on `PATH` beside `swarmfile` (on macOS, `/usr/local/bin`; on Linux, `/usr/bin`). There's no separate download. If `git` says it can't find a helper for `swarmfile`, use `swarmfile git clone`, which finds the helper next to the `swarmfile` binary even when it isn't on `PATH`. Git 2.11 or later is required.

How git uses it

```
git clone swarmfile://<org>/<project>
git remote add origin swarmfile://<org>/<project>
git fetch origin
git push origin main
```

`<project>` may be a bare project name or a project id. A bare name is qualified with the configured organization by `swarmfile git url` / `swarmfile git clone`. Git access is turned on per project — from the web project menu, the Desktop App's **Project settings → Git access**, or `swarmfile git enable` — and a clone of a project that hasn't enabled it is refused with `no_git_view`; the org's owners are nudged once, and `swarmfile git status` shows where a project stands.

What it serves

Shallow clones with `--depth`, fetches, fast-forward pushes (merges land on auto-archived `merged/<sha>` branches), lightweight tag moves and deletes, and chunked resumable pushes above 5,000 paths. Force push is refused by default — a clone can opt in with `git config swarmfile.force-forward true` to translate plain `--force` / `--force-with-lease` into a forward commit (the same end state `swarmfile git force-forward` lands); annotated tags, octopus merges, submodules, end-to-end-encrypted projects and HTTPS remotes are refused, and `--shallow-since` / `--shallow-exclude` aren't supported. The full list, and the version-control model around it, is in [Clone and push with git](#) and [Git and Swarmfile](#).

Credentials

The helper looks for, in order:

1. **SWARMFILE_API_KEY** in the environment. A project API key wins over everything else, including a Desktop App session, and the helper says so when both are present.
2. A **Desktop App session** stored by the engine.
3. A **project API key in git's credential store**, stored with `git credential approve` for the Swarmfile host.

When something goes wrong

A failed helper always prints its reason above git's `fatal: remote helper 'swarmfile' aborted session` line, and writes the same text to `git-remote-last-error.log` in the Swarmfile cache directory (`~/.cache/swarmfile` on macOS and Linux, `%LOCALAPPDATA%\Swarmfile` on Windows). The message names who can fix it; include that log in a bug report.