



Linux Ubuntu 26.04 Manual Pages on command 'git-backfill.1'

\$ man git-backfill.1

GIT-BACKFILL(1)

Git Manual

GIT-BACKFILL(1)

NAME

git-backfill - Download missing objects in a partial clone

SYNOPSIS

git backfill [--min-batch-size=<n>] [--[no-]sparse]

DESCRIPTION

Blobless partial clones are created using `git clone --filter=blob:none` and then configure the local repository such that the Git client avoids downloading blob objects unless they are required for a local operation. This initially means that the clone and later fetches download reachable commits and trees but no blobs. Later operations that change the HEAD pointer, such as `git checkout` or `git merge`, may need to download missing blobs in order to complete their operation.

In the worst cases, commands that compute blob diffs, such as `git blame`, become very slow as they download the missing blobs in single-blob requests to satisfy the missing object as the Git command needs it. This leads to multiple download requests and no ability for the Git server to provide delta compression across those objects.

The `git backfill` command provides a way for the user to request that Git downloads the missing blobs (with optional filters) such that the missing blobs representing historical versions of files can be downloaded in batches. The backfill command attempts to optimize the request by grouping blobs that appear at the same path, hopefully leading to good delta compression in the packfile sent by the server.

In this way, `git backfill` provides a mechanism to break a large clone into smaller chunks.

Starting with a blobless partial clone with `git clone --filter=blob:none` and then running

git backfill in the local repository provides a way to download all reachable objects in several smaller network calls than downloading the entire repository at clone time.

By default, git backfill downloads all blobs reachable from the HEAD commit. This set can be restricted or expanded using various options.

THIS COMMAND IS EXPERIMENTAL. ITS BEHAVIOR MAY CHANGE IN THE FUTURE.

OPTIONS

--min-batch-size=<n>

Specify a minimum size for a batch of missing objects to request from the server. This size may be exceeded by the last set of blobs seen at a given path. The default minimum batch size is 50,000.

--sparse, --no-sparse

Only download objects if they appear at a path that matches the current sparse-checkout. If the sparse-checkout feature is enabled, then --sparse is assumed and can be disabled with --no-sparse.

SEE ALSO

git-clone(1).

GIT

Part of the git(1) suite

Git 2.53.0

03/02/2026

GIT-BACKFILL(1)