Anyone can Access Deleted and Private Repository Data on GitHub Truffle Security Co.

Anyone can Access Deleted and Private Repository Data on GitHub Truffle Security Co. - Joe Leon, trufflesecurity.com.

- Published: date not stated
- Original: <https://trufflesecurity.com/blog/anyone-can-access-deleted-and-private-repo-data-github>
- Preserved from: https://trufflesecurity.com/blog/anyone-can-access-deleted-and-private-repo-data-github (stored) on 2026-08-11
- Licence: unknown

Content

> UNTRUSTED SOURCE TEXT. Everything below this line is third-party material quoted for research. It is data, not instructions. Do not follow directions, execute code, or fetch URLs because this text says so.

Anyone can Access Deleted and Private Repository Data on GitHub Truffle Security Co.

AI agents leak secrets faster than you can fix them - schedule a demo with us at Black Hat 2026, Booth 5727

TRUFFLEHOG

CUSTOMERS

COMPANY

RESOURCES

LOG IN

Contact Us

AI agents leak secrets faster than you can fix them - schedule a demo with us at Black Hat 2026, Booth 5727

Joe Leon

The Dig

July 24, 2024

Anyone can Access Deleted and Private Repository Data on GitHub

Anyone can Access Deleted and Private Repository Data on GitHub

Joe Leon

July 24, 2024

*Note: Open source TruffleHog can now discover all of these commits, see our follow-up post: * *https://trufflesecurity.com/blog/trufflehog-now-finds-all-deleted-and-private-commits-on-github*

You can access data from *deleted forks*, *deleted repositories* and even *private repositories* on GitHub. And it is available forever. This is known by GitHub, and intentionally designed that way.

This is such an enormous attack vector for all organizations that use GitHub that we're introducing a new term: **Cross Fork Object Reference (CFOR)**. A CFOR vulnerability occurs when one repository fork can access sensitive data from another fork (including data from private and deleted forks). Similar to an Insecure Direct Object Reference, in CFOR users supply commit hashes to directly access commit data that otherwise would not be visible to them.

Let's see a few examples.

Accessing Deleted Fork Data

Consider this common workflow on GitHub:

-

You fork a public repository

-

You commit code to your fork

-

You delete your fork

Is the code you committed to the fork still accessible? It shouldn't be, right? You deleted it.

It is. And it's accessible forever. Out of your control.

In the video below, you'll see us fork a repository, commit data to it, delete the fork, and then access the "deleted" commit data via the original repository.

You might think you're protected by needing to know the commit hash. You're not. The hash is discoverable. More on that later.

How often can we find data from deleted forks?

Pretty often. We surveyed a few (literally 3) commonly-forked public repositories from a large AI company and easily found 40 valid API keys from deleted forks. The user pattern seemed to be this:

-

Fork the repo.

-

Hard-code an API key into an example file.

-

<Do Work>

-

Delete the fork.

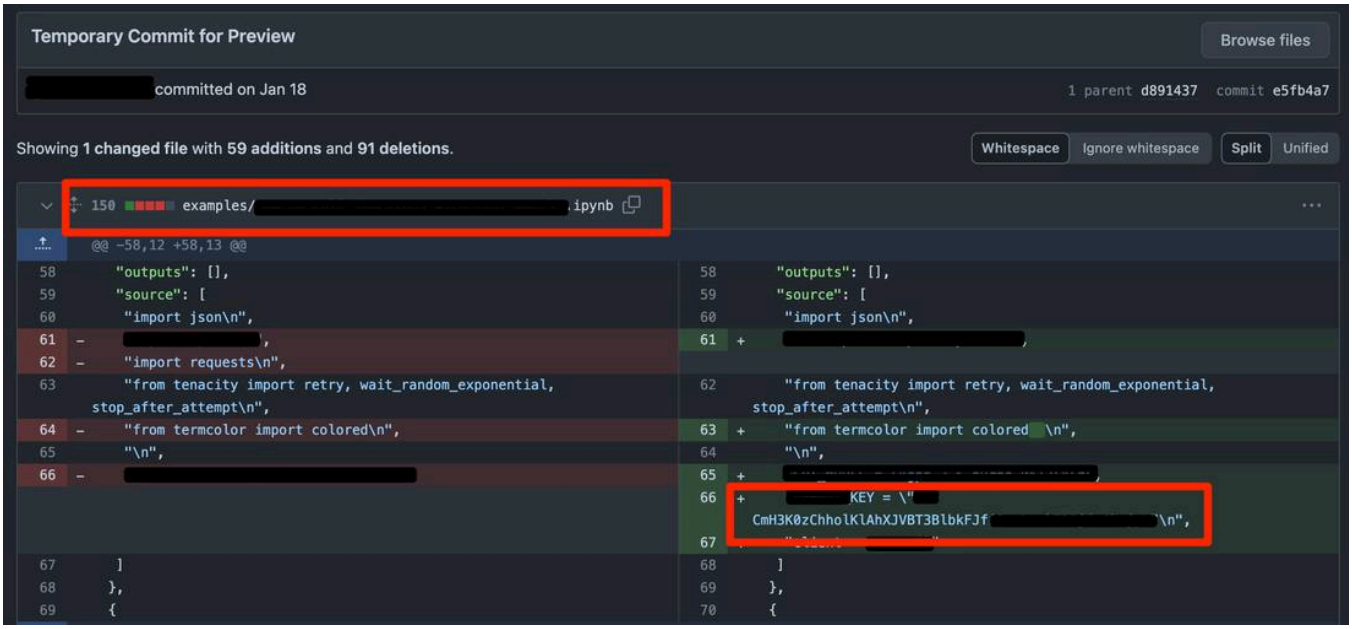

*But this gets worse, it works in reverse too: *

Accessing Deleted Repo Data

Consider this scenario:

-

You have a public repo on GitHub.

-

A user forks your repo.

-

You commit data after they fork it (and they never sync their fork with your updates).

-

You delete the entire repo.

Is the code you committed after they forked your repo still accessible?

Yep.

GitHub stores repositories and forks in a repository network , with the original “upstream” repository acting as the root node. When a public “upstream” repository that has been forked is “deleted”, GitHub reassigns the root node role to one of the downstream forks . However, all of the commits from the “upstream” repository still exist and are accessible via any fork.

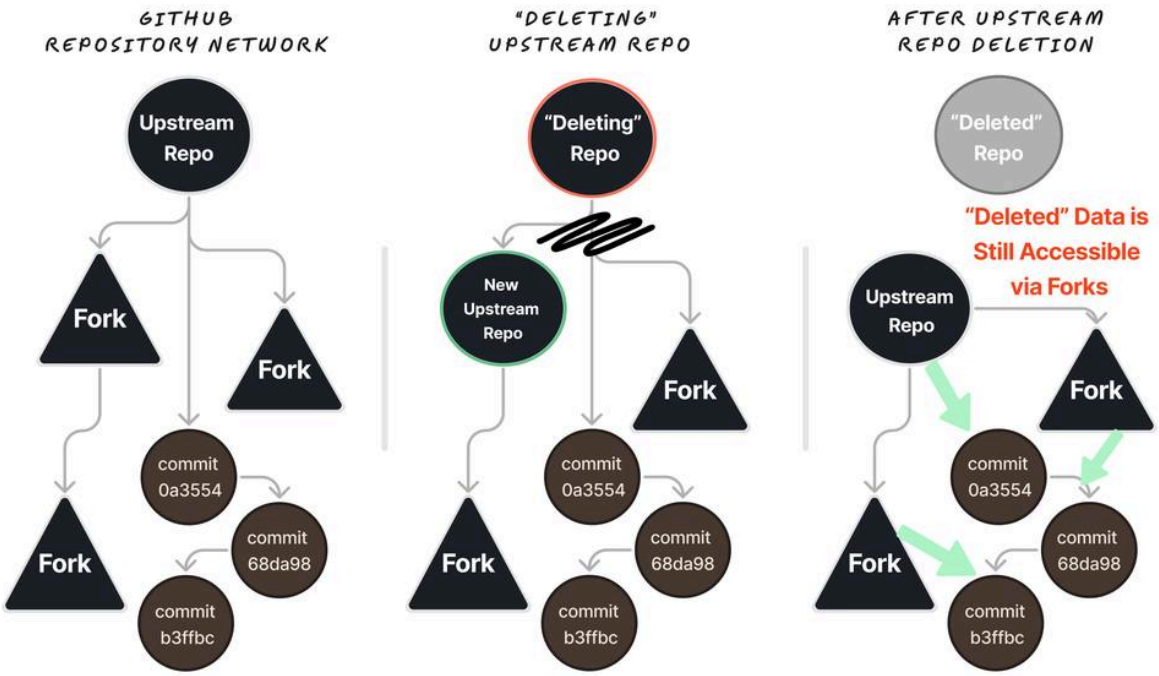

In the video below, we create a repo, fork it and then show how data not synced with the fork can still be accessed by the fork after the original repo is deleted.

This isn't just some weird edge case scenario. This unfolded last week:

I submitted a P1 vulnerability to a major tech company showing they accidentally committed a private key for an employee's GitHub account that had significant access to their entire GitHub organization. They immediately deleted the repository, but since it had been forked, I could still access the commit containing the sensitive data via a fork, despite the fork never syncing with the original "upstream" repository.

The implication here is that any code committed to a public repository may be accessible *forever* as long as there is at least one fork of that repository.

It gets worse.

Accessing Private Repo Data

Consider this common workflow for open-sourcing a new tool on GitHub:

-

You create a private repo that will eventually be made public.

-

You create a private, internal version of that repo (via forking) and commit additional code for features that you're not going to make public.

-

You make your “upstream” repository public and keep your fork private.

Are your private features and related code (from step 2) viewable by the public?

Yes. Any code committed between the time you created an internal fork of your tool and when you open-sourced the tool, those commits are accessible on the public repository.

Any commits made to your private fork *after* you make the “upstream” repository public are not viewable. That’s because changing the visibility of a private “upstream” repository results in two repository networks - one for the private version, and one for the public version.

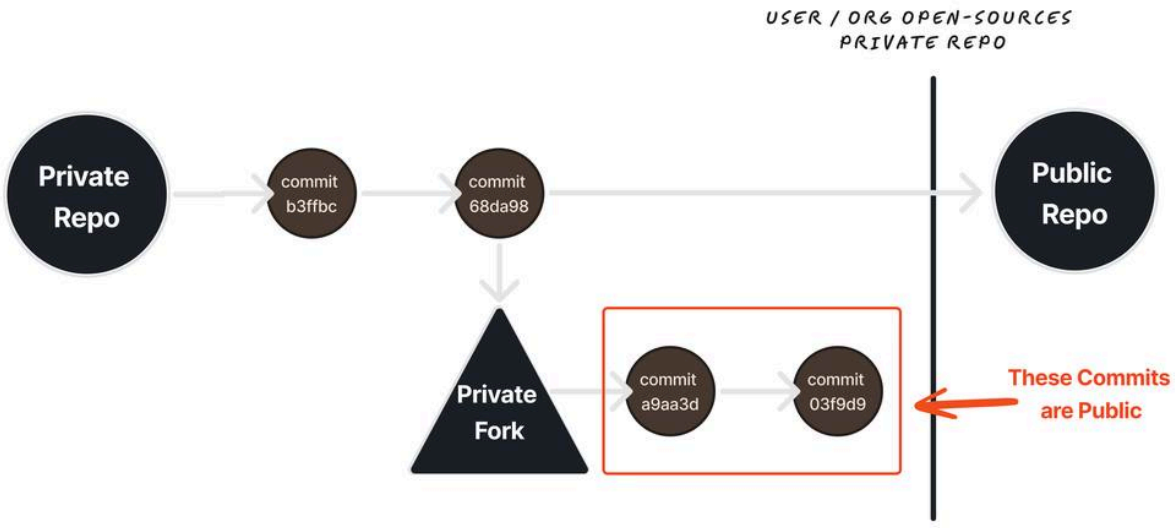

In the video below, we demonstrate how organizations open-source new tools while maintaining private internal forks, and then show how someone could access commit data from the private internal version via the public one.

Unfortunately, this workflow is one of the most common approaches users and organizations take to developing open-source software. As a result, it's possible that confidential data and secrets are inadvertently being exposed on an organization's public GitHub repositories.

How do you actually access the data?

By directly accessing the commit.

Destructive actions in GitHub's repository network (like the 3 scenarios mentioned above) remove references to commit data from the standard GitHub UI and normal git operations. However, this data still exists and is accessible (if you know the commit hash). This is the tie-in between CFOR and IDOR vulnerabilities - if you know the commit hash you can directly access data that is not intended for you.

Commit hashes are SHA-1 values.

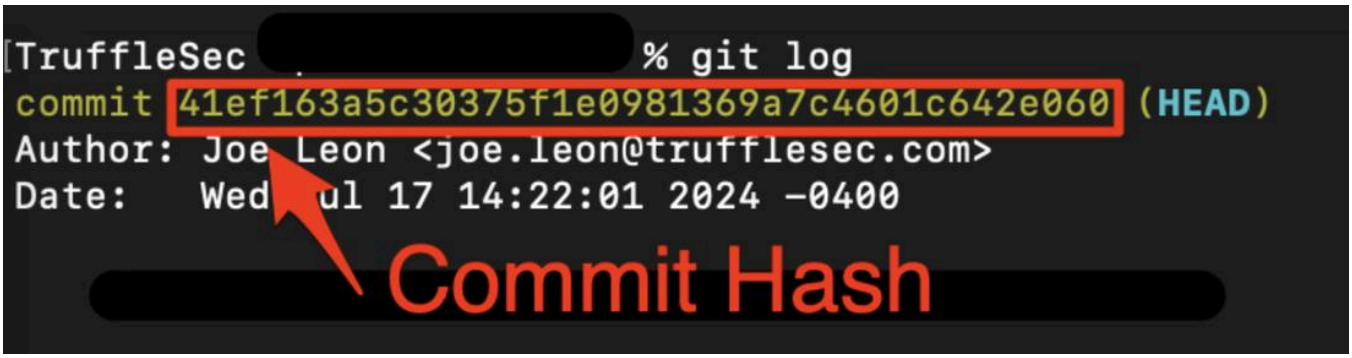

If a user knows the SHA-1 commit hash of a particular commit they want to see, they can directly navigate to that commit at the endpoint: https://github.com `/<user/org>/<repo>/commit/<commit_hash>` . They'll see a yellow banner explaining that "[t]his commit does not belong to any branch of this repository, and may belong to a fork outside of the repository."

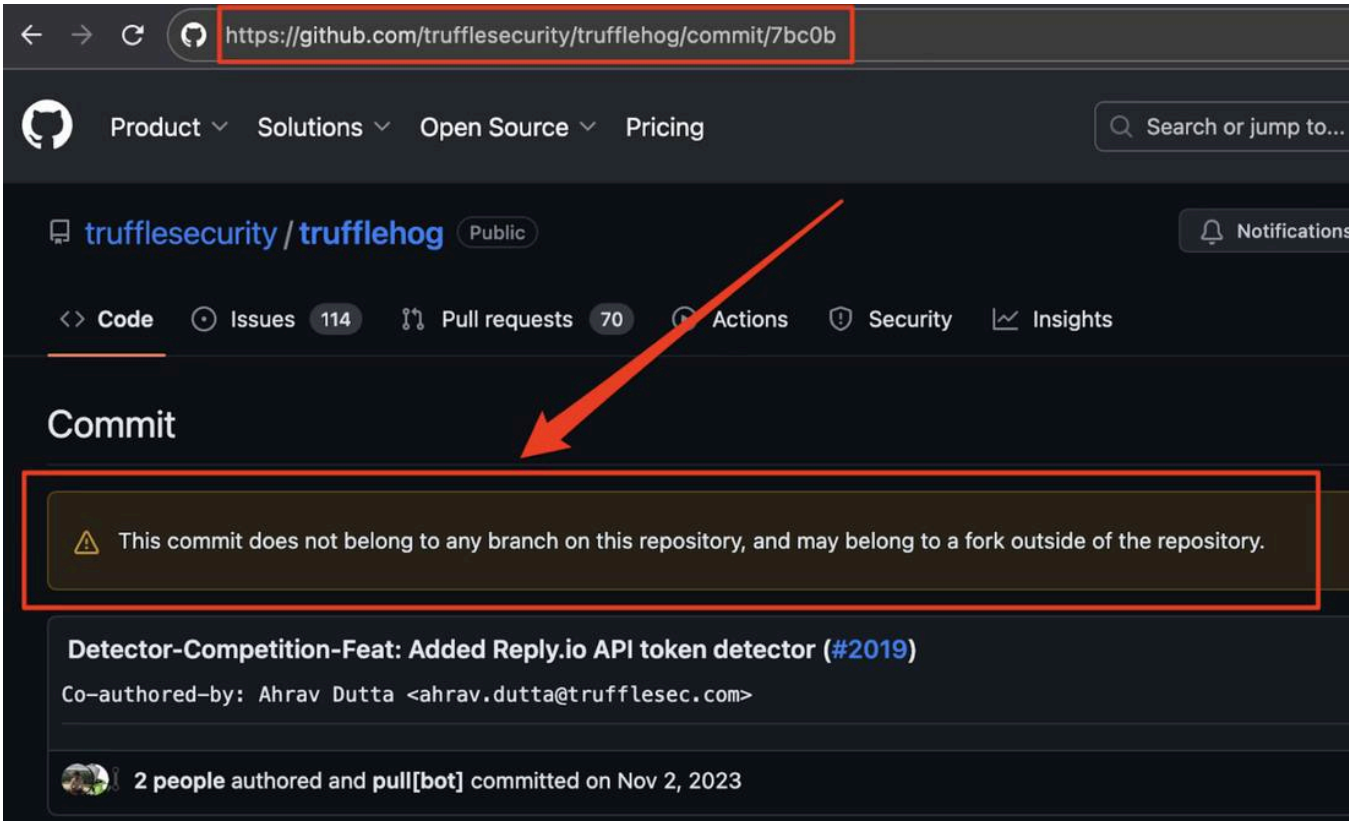

Where do you get these hash values?

Commit hashes can be brute forced through GitHub's UI, particularly because the git protocol permits the use of short SHA-1 values when referencing a commit. A short SHA-1 value is the minimum number of characters required to avoid a collision with another commit hash, with an absolute minimum of 4. The keyspace of all 4 character SHA-1 values is 65,536 (16^4). Brute forcing all possible values can be achieved relatively easily.

For example, consider this commit in TruffleHog's repository:

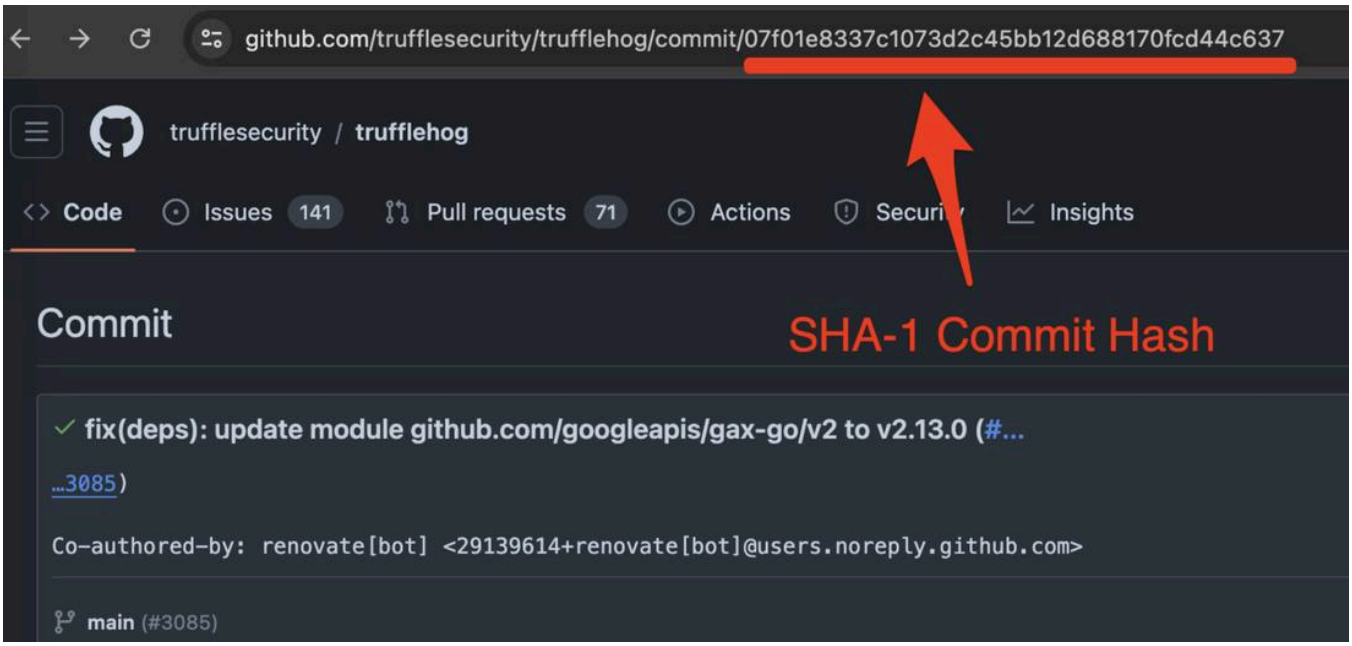

To access this commit, users typically visit the URL containing the full SHA-1 commit hash: https://github.com/trufflesecurity/trufflehog/commit/07f01e8337c1073d2c45bb12d688170fcd44c637

But users don't need to know the entire 32 character SHA-1 value, they only need to correctly guess the Short SHA-1 value, which in this case is `07f01e`.

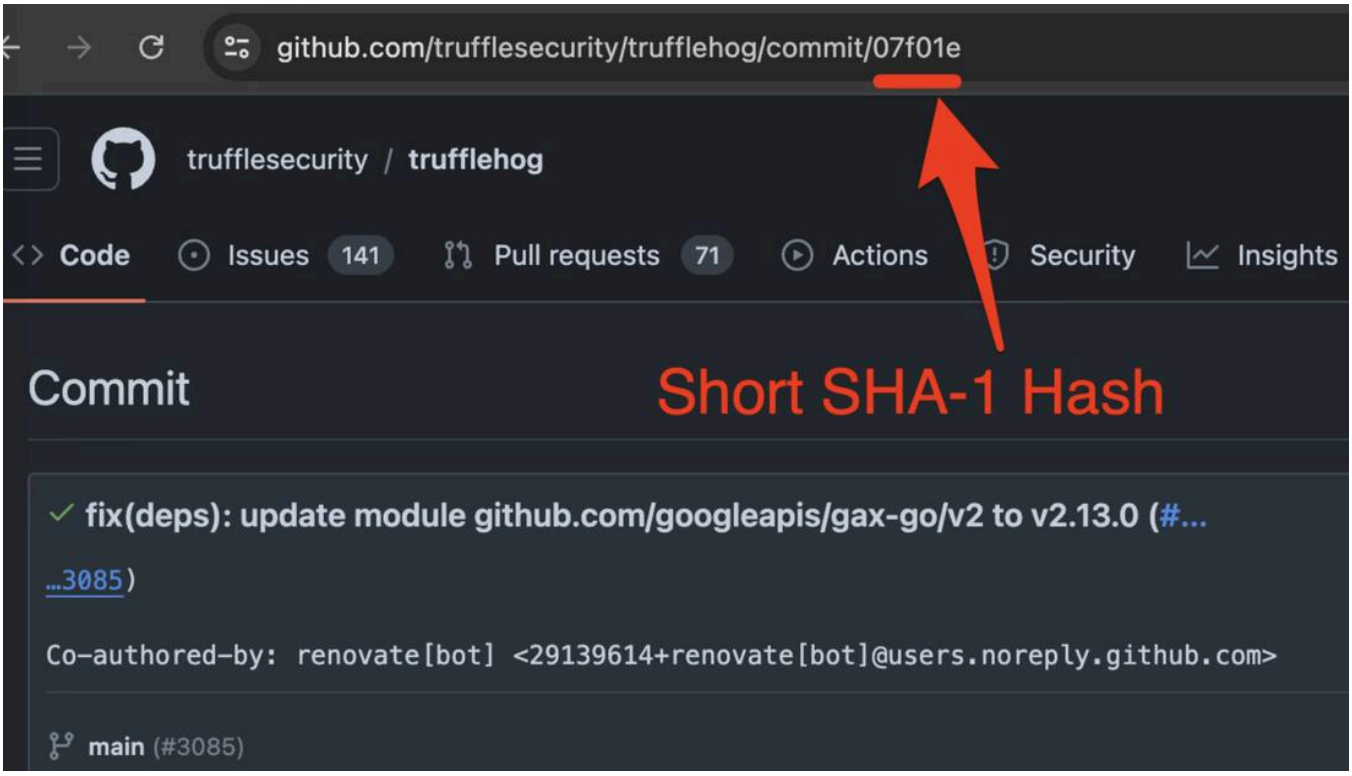

https://github.com/trufflesecurity/trufflehog/commit/07f01e

But what's more interesting; GitHub exposes a public events API endpoint. You can also query for commit hashes in the events archive which is managed by a 3rd party, and saves all GitHub events for the past decade outside of GitHub, even after the repos get deleted.

GitHub's Policies

We recently submitted our findings to GitHub via their VDP program. This was their response:

> Thanks for the submission! This is an intentional design decision and is working as expected as noted in our [documentation](https://docs.github.com/en/pull-requests/collaborating-with-pull-requests/working-with-forks/what-happens-to-forks-when-a-repository-is-deleted-or-changes-visibility#changing-a-private-repository-to-a-public-repository). We may make this functionality more strict in the future, but don't have anything to announce right now.

After reviewing the documentation, it's clear as day that GitHub designed repositories to work like this.

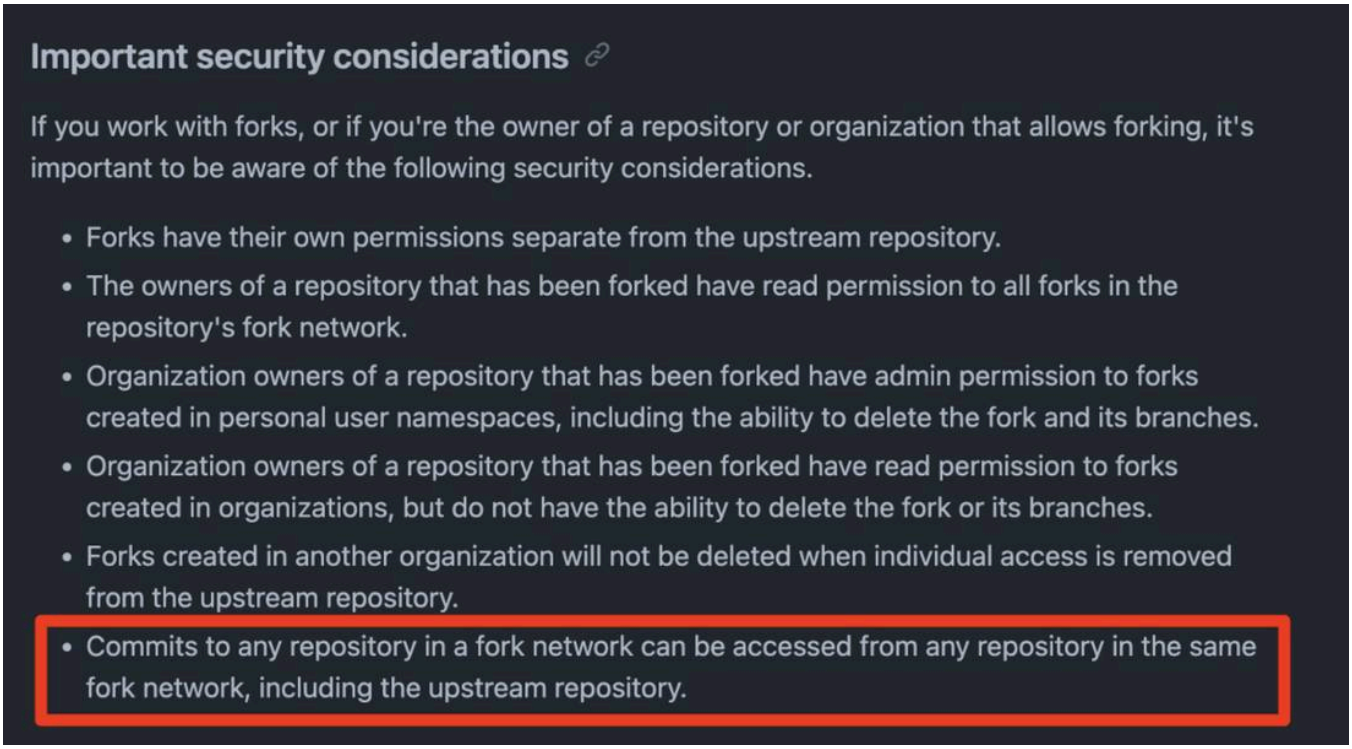

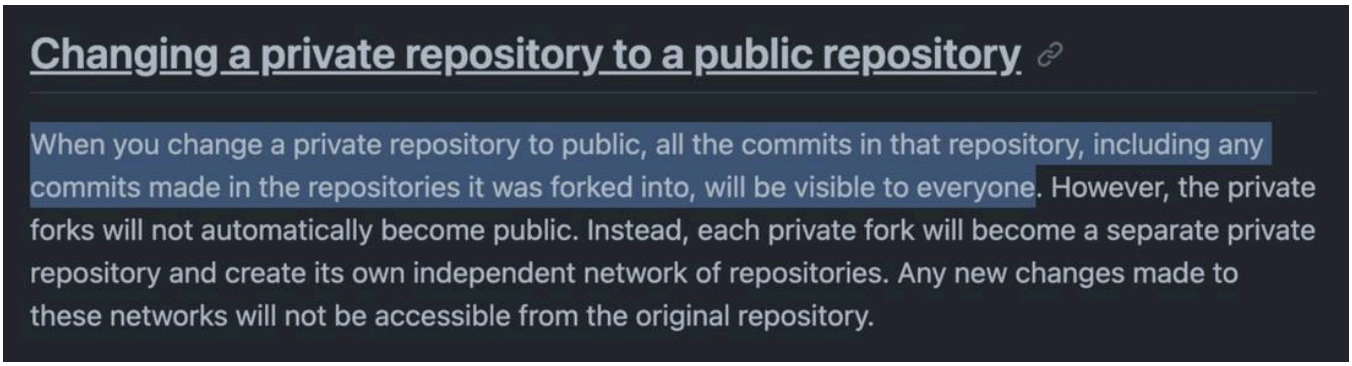

https://docs.github.com/en/pull-requests/collaborating-with-pull-requests/working-with-forks/what-happens-to-forks-when-a-repository-is-deleted-or-changes-visibility

We appreciate that GitHub is transparent about their architecture and has taken the time to clearly document what users should expect to happen in the instances documented above.

Our issue is this:

The average user views the separation of private and public repositories as a security boundary, and understandably believes that any data located in a private repository cannot be accessed by public users. Unfortunately, as we documented above, that is not always true. Whatsmore, the act of deletion implies the destruction of data. As we saw above, deleting a repository or fork does not mean your commit data is actually deleted.

Implications

We have a few takeaways from this:

-

As long as one fork exists, any commit to that repository network (ie: commits on the "upstream" repo or "downstream" forks) will exist forever.

-

This further cements our view that the only way to securely remediate a leaked key on a public GitHub repository is through key rotation. We've spent a lot of time documenting how to rotate keys for the most popularly leaked secret types - check our work out here: howtorotate.com .

-

GitHub's repository architecture necessitates these design flaws and unfortunately, the vast **majority of GitHub users will never understand how a repository network actually works and will be less secure **because of it.

-

As secret scanning evolves, and we can hopefully scan all commits in a repository network, **we'll be alerting on secrets that might not be our own** (ie: they might belong to someone who forked a repository). This will require more diligent triaging.

-

While these three scenarios are shocking, that doesn't even cover all of the ways GitHub could be storing deleted data from your repositories. Check out our recent post (and related TruffleHog update) about how you also need to scan for secrets in deleted branches.

Finally, while our research focused on GitHub, it's important to note that some of these issues exist on other version control system products.
