Everything I Know About the XZ Backdoor

By Evan Boehs

Published: 2024-03-29 · Archived: 2026-04-05 15:04:36 UTC

state

evergreen

in

blog

tags

- open-source

date

3/29/2024

This publication was last updated at 12:49 PM EST on April 8th

Recently, a backdoor was discovered in XZ, a popular library for lossless data compression. Initial research efforts were predominantly concentrated on unpacking the well-disguised attack vector, while the social aspects of the attack received only murmurings. To investigate this attack, I never read a line of code. Instead, I spent dozens of hours pouring over hundreds of discussion threads and mailing lists. I've concluded that while the events that unfolded are undoubtedly a tragedy, it was one of our own making — one of an ugly reality that quietly plays out every day. Our attacker was not naïve. They meticulously examined the culture of open source and then pounced on its norms, twisting them strategically to their benefit. If we continue to focus purely on technicalities, nothing will ever change. This story will play out time and time again, like it already has for decades. We must stop focusing on fuzzing and static analysis and instead turn our attention to the human costs of open source. We must learn. With this in mind, I present a timeline of the attack — one from a new perspective that's too often ignored.

2021

JiaT75 (Jia Tan) creates their GitHub account.

The first commits they make are not to xz, but they are deeply suspicious. Specifically, they open a PR in libarchive: Added error text to warning when untaring with bsdtar. This commit does a little more than it says. It replaces `safe_fprint` with an unsafe variant, potentially introducing another vulnerability. The code was merged without any discussion, and ~~lives on to this day~~ (patched).

2022

In April 2022, Jia Tan submitted a patch via a mailing list. The contents of the patch are not relevant, but the events that follow are. A new persona — *Jigar Kumar* — enters, and begins pressuring for this patch to be

merged.

Soon after, *Jigar Kumar* begins pressuring *Lasse Collin* to add another maintainer to XZ. In the fallout, there is much to learn about mental health in open source.

Three days after the emails pressuring *Lasse Collin* to add another maintainer, JiaT75 makes their first commit to xz: Tests: Created tests for hardware functions.. Since this commit, they become a regular contributor to xz (they are currently the second most active). It's unclear exactly when they became trusted in this repository.

Jigar Kumar is never seen again. Another account — Dennis Ens also participates in pressure, with a similar name+number formatted email. This account is also never seen outside of xz discussion, and neither have any associated accounts that have been discovered.

Glyph @glyph@mastodon.social

@eb I really hope that this causes an industry-wide reckoning with the common practice of letting your entire goddamn product rest on the shoulders of one overworked person having a slow mental health crisis without financially or operationally supporting them whatsoever. I want everyone who has an open source dependency to read this message https://www.mail-archive.com/xz-devel@tukaani.org/msg00567.html

Mar 29, 2024, 20:43 536 retoots

2023

JiaT75 merges their first commit on Jan 7, 2023[1], which gives us a good indication of when they fully gain trust.

In March, the primary contact email in Google's oss-fuzz is updated to be Jia's, instead of *Lasse Collin*.

Testing infrastructure that will be used in this exploit is committed. Despite *Lasse Collin* being attributed as the author for this, *Jia Tan* committed it, and it was originally written by *Hans Jansen* in June:

- Commit: liblzma: Add ifunc implementation to crc64_fast.c
- PR: Replaced crc64_fast constructor with ifunc by hansjans162

Hans Jansen's account was seemingly made specifically to create this pull request. There is very little activity before and after. They will later push for the compromised version of XZ to be included in Debian.

In July, a PR was opened in oss-fuzz to disable ifunc for fuzzing builds, due to issues introduced by the changes above. This appears to be deliberate to mask the malicious changes that will be introduced soon. Also, JiaT75 opened an issue about a warning in clang that, while indeed incorrect, drew attention to ifuncs.

2024

A pull request for Google's oss-fuzz is opened that changes the URL for the project from tukaani.org/xz/ to xz.tukaani.org/xz-utils/. tukaani.org is hosted at `5.44.245.25` in Finland, at this hosting company. The `xz` subdomain, meanwhile, points to GitHub pages. This furthers the amount of control Jia has over the project.

A commit containing the final steps required to execute this backdoor is added to the repository:

- Tests: Add a few test files
- Tests: Update two test files

The discovery

An email is sent to the oss-security mailing list: backdoor in upstream xz/liblzma leading to ssh server compromise, announcing this discovery, and doing it's best to explain the exploit chain.

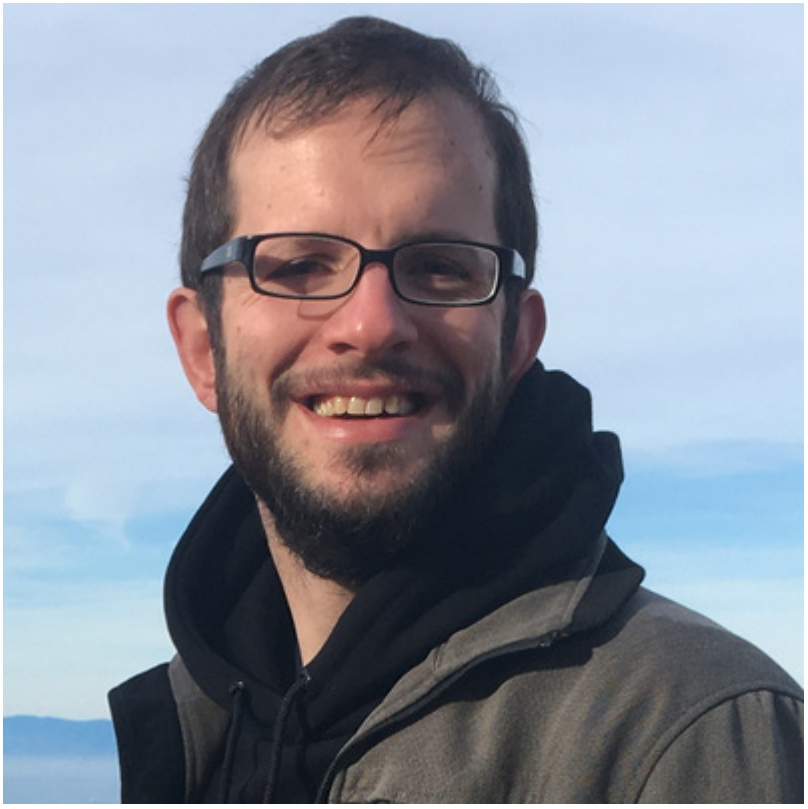

I was doing some micro-benchmarking at the time, needed to quiesce the system to reduce noise. Saw sshd processes were using a surprising amount of CPU, despite immediately failing because of wrong usernames etc. Profiled sshd, showing lots of cpu time in liblzma, with perf unable to attribute it to a symbol. Got suspicious. Recalled that I had seen an odd valgrind complaint in automated testing of postgres, a few weeks earlier, after package updates.

Really required a lot of coincidences.

Mar 29, 2024, 18:32 769 retoots

A gist has been published with a great high-level technical overview and a “what you need to know”

In addition to the gist and the email above, several analysis attempts have begun emerging:

- xz/liblzma: Bash-stage Obfuscation Explained
- “It’s RCE, not auth bypass”
- [WIP] XZ Backdoor Analysis and symbol mapping
- Infographic
- xzbot: notes, honeypot, and exploit demo for the xz backdoor
- research!rsc: The xz attack shell script

A sudden push for inclusion

A request for the vulnerable version to be included in Debian is opened by Hans:

- #1067708 - xz-utils: New upstream version available

This request was opened the same week Hans’ Debian GitLab account was created. The account created a few similar “update” requests in various low-traffic repositories to build credibility, after asking for this one.

Several other, suspicious, anonymous name+number accounts with little former activity also push for its inclusion, including *misoeater91* and *krygorin4545*. *krygorin4545*’s PGP key was made 2 days before joining the discussion.

> Also seeing this bug. Extra valgrind output causes some failed tests for me. Looks like the new version will resolve it. Would like this new version so I can continue work.

> I noticed this last week and almost made a Valgrind bug. Glad to see it being fixed.
> Thanks Hans!

The Valgrind bugs mentioned were *introduced* by this malicious injection, as noted in the email to OSS-Security:

> Subsequently the injected code (more about that below) caused valgrind errors and crashes in some configurations, due to the stack layout differing from what the backdoor was expecting. These issues were attempted to be worked around in 5.6.1:

A pull request to a go library by a 1Password employee is opened asking to upgrade the library to the vulnerable version, however, it was all unfortunate timing. 1Password reached out by email referring me to this comment, and everything seems to check out.

A Fedora contributor states that *Jia* was pushing for its inclusion in Fedora as it contains “great new features”

Jia Tan also attempted to get it into Ubuntu days before the beta freeze.

A few hours after all this came out, GitHub suspended JiaT75’s account. Thanks? They also banned the repository, meaning people can no longer audit the changes made to it without resorting to mirrors. Immensely helpful, GitHub. They also suspended *Lasse Collin*’s account, which is completely disgraceful.

Lasse has begun reverting changes introduced by Jia, including one that added a sneaky period to disable the sandbox. They also have published a FAQ that begins to explain the situation: XZ Utils backdoor

OSINT

Various people have reached out to me regarding discoveries about the identity of Jia. Some of this has been incorporated in the timeline, but other stuff is "timeless" so I'm putting it here:

IRC

I received an email that clarified a few points and provided new insight into the situation.

"Jia Tan" was present on the #tukaani IRC channel on Libera.Chat. A /whois revealed their connecting IP and activity on March 29th.

```
 [libera] -!- jiatan [~jiatan@185.128.24.163]
[libera] -!-  was      : Jia Tan
[libera] -!-  hostname : 185.128.24.163
[libera] -!-  account  : jiatan
[libera] -!-  server   : tungsten.libera.chat [Fri Mar 29 14:47:40 2024]
[libera] -!- End of WHOWAS
```

Running a Nmap on the IP shows a lot of open ports, which probably indicates a proxy, hosting provider, or something of the sort. The IP is from Singapore.

Further research shows that this IP belongs to Witopia VPN, so it's not entirely indicative of a region. Given the timezone, however, I feel like proximity becomes plausible.

Important notes on LinkedIn

I have received a few emails alerting me to a LinkedIn of somebody named *Jia Tan*[2]. Their bio boasts of *large-scale vulnerability management*. They claim to live in California. Is this our man? The commits on JiaT75's GitHub are set to +0800, which would not indicate presence in California. UTC-0800 would be California. Most of the commits were made between UTC 12-17, which is awfully early for California. In my opinion, there is no sufficient evidence that the LinkedIn being discussed is our man. I think identity theft is more likely, but I am of course open to more evidence.

Apr 7 Update: Subsequently, I've received a lot of people sending me other LinkedIn accounts, theorizing about what his name could mean, etc., but I don't think a bad actor would use their real name.

Discoveries in the Git logs

I received an email from Minhu Wang who investigated the Git log, and found one instance where Jia's username was different:

```
 $ git shortlog --summary --numbered --email | grep jiat0218@gmail.com
273 Jia Tan <jiat0218@gmail.com>
2 jiat75 <jiat0218@gmail.com>
1 Jia Cheong Tan <jiat0218@gmail.com>
```

They found this particularly interesting as `Cheong` is new information

Furthermore, an independent analysis of commit timings concludes that the perpetrator worked “Office Hours” in a UTC+02/03 timezone. It’s particularly notable that they worked through the Lunar New Year, and did not work on some notable Eastern European holidays, including Christmas and New Year. I have, however, been presented with a differing view, which you can read here.

Ry Jones used gharchive to extract Jia’s entire git activity, and he uploaded it to a repository, viewable here: jiat75-logs. Jia’s GitHub username previously featured Cheong, but this has since been removed.

👟 Footnotes

1. Thanks @joeyh@hachyderm.io ↩

2. I was also alerted to discussions of this on Gab, which should tell you what you need to know. ↩

🌱 What links here?

- Bountysource Stole at Least $21,000 From Open Source Developers
- How did the blog do in 2024?
