securelist.com /xz-backdoor-story-part-2-social-engineering/112476/

Assessing the Y, and How, of the XZ Utils incident

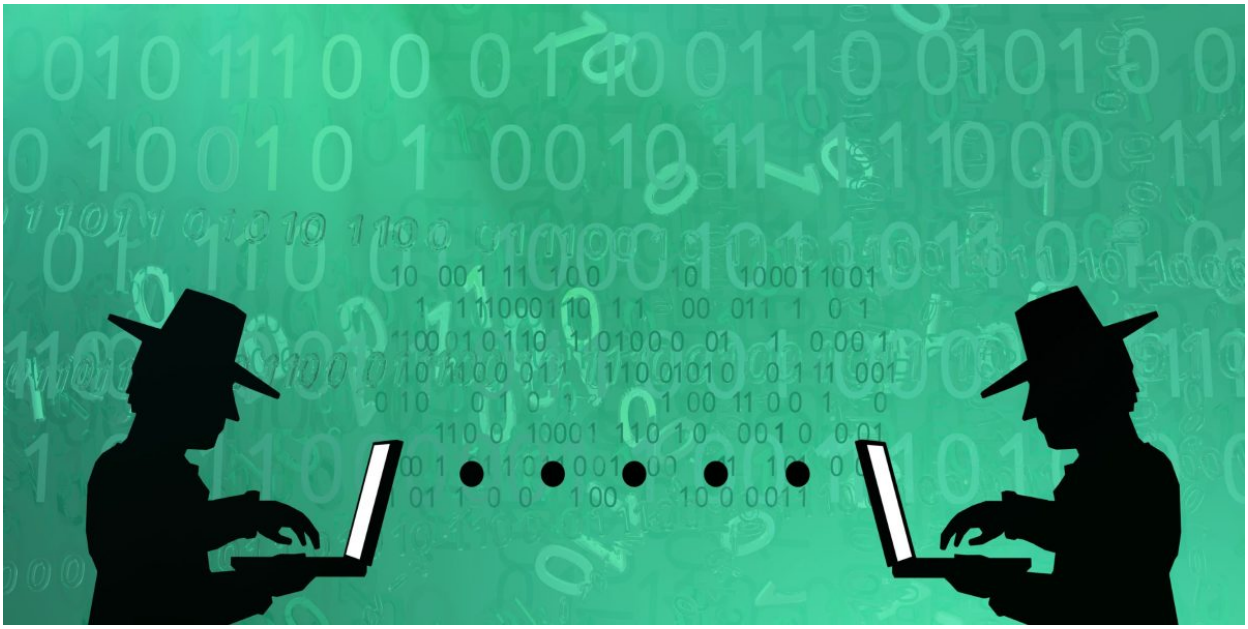

Authors

- Expert GReAT

Social engineering for open-source supply chain attack profit

High-end APT groups perform highly interesting social engineering campaigns in order to penetrate well-protected targets. For example, carefully constructed forum responses on precision targeted accounts and follow-up “out-of-band” interactions regarding underground rail system simulator software helped deliver Green Lambert implants in the Middle East. And, in what seems to be a learned approach, the XZ Utils project penetration was likely a patient, multi-year approach, both planned in advance but somewhat clumsily executed.

This recently exposed offensive effort slowly introduced a small cast of remote characters, communications, and malicious code to the more than decade old open-source project XZ Utils and its maintainer, Lasse Collin. The backdoor code was inserted in February and March 2024, mostly by Jia Cheong Tan, likely a fictitious identity. The end goal was to covertly implement an exclusive use backdoor in sshd by targeting the XZ Utils build process, and push the backdoored code to the major Linux distributions as a part of a large-scale supply chain attack.

While this highly targeted and interactive social engineering approach might not be completely novel, it is extraordinary. Also extraordinary is the stunningly subtle insertion of malicious code leveraging the build process in plain sight. This build process focus during a major supply chain attack is comparable only to the CozyDuke/DarkHalo/APT29/NOBELIUM Solarwinds compromise and the SUNSPOT implant’s

cunning and persistent presence – its monitoring capability for the execution of a Solarwinds build, and its malicious code insertion during any Solarwinds build execution. Only this time, it's human involvement in the build process.

It's notable that one of the key differentiators of the Solarwinds incident from prior supply chain attacks was the adversary's covert, prolonged access to the source/development environment. In this XZ Utils incident, this prolonged access was obtained via social engineering and extended with fictitious human identity interactions in plain sight.

One of the best publicly available chronological timelines on the social engineering side of the XZ Utils incident is posted by Russ Cox, currently a Google researcher. It's highly recommended reading. Notably, Cox writes: "This post is a detailed timeline that I have constructed of the social engineering aspect of the attack, which appears to date back to late 2021."

A Singaporean guy, an Indian guy, and a German guy walk into a bar…

Three identities pressure XZ Utils creator and maintainer Lasse Collin in summer 2022 to provoke an open-source code project handover: Jia Tan/Jia Cheong Tan, Dennis Ens, and Jigar Kumar. These identities are made up of a GitHub account, three free email accounts with similar name schemes, an IRC and Ubuntu One account, email communications on XZ Utils developer mailing lists and downstream maintainers, and code. Their goal was to grant full access to XZ Utils source code to Jia Tan and subtly introduce malicious code into XZ Utils – the identities even interact with one another on mail threads, complaining about the need to replace Lasse Collin as the XZ Utils maintainer.

Note that the geographic dispersion of fictitious identities is a bit forced here, perhaps to dispel hints of coordination: Singaporean or Malaysian (possibly of a Hokkien dialect), northern European, and Indian. Misspellings and grammar mistakes are similar across the three identities' communications. The "Jia Tan" identity seems a bit forced as well – the only public geolocation data is a Singaporean VPN exit node that the identity may have used on March 29 to access the XZ Utils Libera IRC chat. If constructing a fictitious identity, using that particular exit node would definitely be a selected resource.

```
[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
```

Our pDNS confirms this IP as a Witopia VPN exit. While we might expect a "jiat75" or "jiatan018" username for the "Jia Tan" Libera IRC account, this one in the screenshot above may have been used on

March 29, 2024 by the “JiaT75” actor.

Host	IP	Country	Firstseen	Lastseen	Countseen	Tags Host IP
vpn.sin.witopia.net	185.128.24.163	SG	2019-02-20 07:17:18	2019-02-26 06:47:40	1	VPN
vpn.singapore.witopia.net	185.128.24.163	SG	2020-03-27 15:43:42	2023-07-15 10:29:13	878	VPN

One additional identity, Hans Jansen, introduced a June 2023 performance optimization into the XZ Utils source, committed by Collin, and later leveraged by jiaT75’s backdoor code. Jia Tan gleefully accepted the proposed IFUNC additions: “Thanks for the PR and the helpful links! Overall this seems like a nice improvement to our function-picking strategy for CRC64. It will likely be useful when we implement CRC32 CLMUL too :)”.

This pull request is the Jansen identity’s only interaction with the XZ Utils project itself. And, unlike the other two identities, the Jansen account is not used to pressure Collin to turn over XZ Utils maintenance. Instead, the Hans Jansen identity provided the code and then disappeared. Nine months later, following the backdoor code insertion, Jansen urged a major Linux vendor in the supply chain to incorporate the backdoored XZ Utils code in their distribution. The identity resurfaced on a Debian bug report on March 24, 2024, creating an opportunity to generate urgency in including the backdoored code in the Debian distribution.

Jia Tan Identity and Activity

The Jia Cheong Tan (JiaT75) GitHub account, eventually promoted to co-maintainer of XZ Utils, which inserted the malicious backdoor code, was created January 26, 2021. JiaT75 was not exclusively involved in XZ Utils, having authored over 500 patches to multiple GitHub projects going back to early 2022.

- oss-fuzz
- cpp-docs
- wasmtime
- xz

These innocuous patches helped to build the identity of JiaT75 as a legitimate open source contributor and potential maintainer for the XZ Utils project. The patch efforts helped to establish a relationship with Lasse Collin as well.

The first JiaT75 code contribution to XZ Utils occurred on October 29, 2021. It was sent to the xz-devel mailing list. It was a very simple editor config file introduction. Following this initial innocuous addition, over the next two years, JiaT75 authored hundreds of changes for the XZ project.

Yes, JiaT75 contributed code on both weekends and what appear to be workdays. However, an interesting anomaly is that the 2024 malicious commits occur out of sync with many previous commits. A Huntress researcher going by the alias “Alden” posted a visualization of the malicious Jia Tan commits to XZ Utils. JiaT75 commits the malicious code completely out of sync with prior work times on Feb 23–26, and March 8 and 9, 2024.

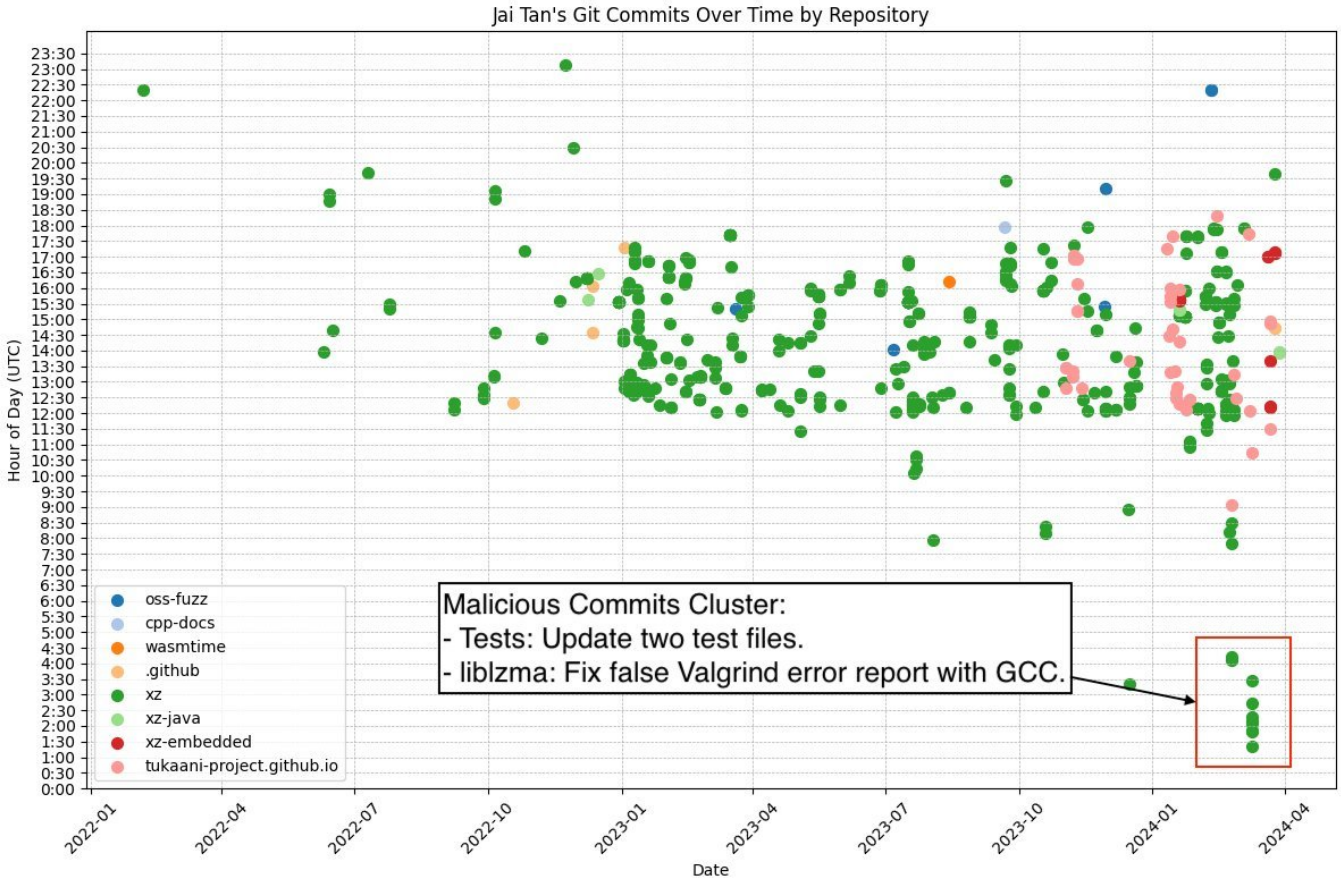

The time differences for the malicious commits is noticeable. What might this anomaly suggest? We speculate on several possibilities:

- the JiaT75 account was used by a second party to insert the malicious code, either known or unknown to the individual contributor.
- the JiaT75 individual contributor was rushed to commit the malicious backdoor code.
- the JiaT75 account was run by a team of individuals and one part of the team needed to work without interruption outside of the usual constructed work day.

Especially devious is the manner in which the obfuscated backdoor code is introduced in multiple separate pieces by JiaT75. Even though it was open-source, the bulk of the backdoor does not show up in the XZ source-code tree, is not human readable, and was not recognized.

Summer 2022 Pressure to Add a Maintainer

Multiple identities of interest pressured Lasse Collin to add a maintainer over the summer of 2022. The intensity of pressure on Collin varies per account, but they all create opportunities to pressure Collin and interact.

Name	GitHub Account	Email	Creation
Jia Tan/Jia Cheong Tan	JiaT75	jiat0218@gmail.com	January 26, 2021
Dennis Ens	–	dennis3ns@gmail.com	–
Jigar Kumar	–	jigarkumar17@protonmail.com	–

If we take the first interaction on the xz-devel mailing list as the start of the campaign, Jia Tan sent a superficial code patch on September 29, 2021. This timestamp is eight months after the github account creation date. This initial contribution is harmless, but establishes this identity within the open-source project.

A year later, Jigar Kumar pressured Lasse Collin to hand over access to Jia Tan over the spring and summer of 2022 in six chiding comments over two different threads.

Wed, 27 Apr 2022 11:42:57 -0700 Re: [xz-devel] [PATCH] String to filter and filter to string
Your efforts are good but based on the slow release schedule it will unfortunatly be years until the community actually gets this quality of life feature.

Thu, 28 Apr 2022 10:10:48 -0700 Re: [xz-devel] [PATCH] String to filter and filter to string
Patches spend years on this mailing list. 5.2.0 release was 7 years ago. There is no reason to think anything is coming soon.

Fri, 27 May 2022 10:49:47 -0700 Re: [xz-devel] [PATCH] String to filter and filter to string
Over 1 month and no closer to being merged. Not a suprise.

Tue, 07 Jun 2022 09:00:18 -0700 Re: [xz-devel] XZ for Java
Progress will not happen until there is new maintainer. XZ for C has sparse commit log too. Dennis you are better off waiting until new maintainer happens or fork yourself. Submitting patches here has no purpose these days. The current maintainer lost interest or doesn't care to maintain anymore. It is sad to see for a repo like this.

Tue, 14 Jun 2022 11:16:07 -0700 Re: [xz-devel] XZ for Java
With your current rate, I very doubt to see 5.4.0 release this year. The only progress since april has been small changes to test code. You ignore the many patches bit rotting away on this mailing list. Right now you choke your repo. Why wait until 5.4.0 to change maintainer? Why delay what your repo needs?

Wed, 22 Jun 2022 10:05:06 -0700 Re: [xz-devel] [PATCH] String to filter and filter to string
"Is there any progress on this? Jia I see you have recent commits. Why can't you commit this yourself?"

The Dennis Ens identity sets up a thread of their own, and follows up by pressuring maintainer Collin in one particularly forceful and obnoxious message to the list. The identity leverages a personal vulnerability that Collin shared on this thread. The Jigar Kumar identity responds twice to this thread, bitterly complaining about the maintainer: "Dennis you are better off waiting until new maintainer happens or fork yourself."

Thu, 19 May 2022 12:26:03 -0700 XZ for Java
Is XZ for Java still maintained? I asked a question here a week ago and have not heard back. When I view the git log I can see it has not updated in over a year. I am looking for things like multithreaded encoding / decoding and a few updates that Brett Okken had submitted (but are still waiting for merge). Should I add these things to only my local version, or is there a plan for these things in the future?

Tue, 21 Jun 2022 13:24:47 -0700 Re: [xz-devel] XZ for Java
I am sorry about your mental health issues, but its important to be aware of your own limits. I get that this is a hobby project for all contributors, but the community desires more. Why not pass on maintainership for XZ for C so you can give XZ for Java more

attention? Or pass on XZ for Java to someone else to focus on XZ for C? Trying to maintain both means that neither are maintained well.

Reflecting on these data points still leads us to shaky ground. Until more details are publicized, we are left with speculation:

- In a three-year project, a small team successfully penetrated the XZ Utils codebase with a slow and low-pressure campaign. They manipulated the introduction of a malicious actor into the trusted position of code co-maintainer. They then initiated and attempted to speed up the process of distributing malicious code targeting sshd to major vendor Linux distributions
- In a three-year project, an individual successfully penetrated the XZ Utils codebase with a slow and low-pressure campaign. The one individual managed several identities to manipulate their own introduction into the trusted position of open source co-maintainer. They then initiated and attempted to speed up the process of distributing malicious code targeting sshd to major vendor Linux distributions
- In an extremely short timeframe in early 2024, a small team successfully manipulated an individual (Jia Tan) that legitimately earned access to an interesting open-source project as code maintainer. Two other individuals (Jigar Kumar, Dennis Ens) may have coincidentally complained and pressured Collin to hand over the maintainer role. That leveraged individual began inserting malicious code into the project over the course of a couple of weeks.

Spring 2024 Pressure to Import Backdoored Code to Debian

Several identities attempted to pressure Debian maintainers to import the backdoored upstream XZ Utils code to their distribution in March 2024. The Hans Jansen identity created a Debian report log on March 25, 2024 to raise urgency to include the backdoored code: “Dear mentors, I am looking for a sponsor for my package “xz-utils”.”

Name	**Email address**
Hans Jansen	hansjansen162@outlook.com
krygorin4545	krygorin4545@proton.me
misoeater91@tutamail.com	misoeater91@tutamail.com

The thread was responded to within a day by additional identities using the email address scheme name-number@freeservice[.]com:

Date: Tue, 26 Mar 2024 19:27:47 +0000 From: krygorin4545 <krygorin4545@proton.me>
Subject: Re: RFS: xz-utils/5.6.1-0.1 [NMU] — XZ-format compression utilities
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
Date: Tue, 26 Mar 2024 22:50:54 +0100 (CET) From: misoeater91@tutamail.com
Subject: Re: RFS: xz-utils/5.6.1-0.1 [NMU] — XZ-format compression
I noticed this last week and almost made a valgrind bug. Glad to see it being fixed. Thanks Hans!

The code changes received pushback from Debian contributors:

Date: Tue, 26 Mar 2024 22:11:19 +0000 (UTC) From: Thorsten Glaser <tg@debian.org>

Subject: new upstream versions as NMU vs. xz maintenance

Very much *not* a fan of NMUs doing large changes such as
new upstream versions.But this does give us the question, what's up with the
maintenance of xz-utils? Same as with the lack of security
uploads of git, which you also maintain, are you active?

Are you well?

To which one of these likely sock puppet accounts almost immediately responded, in order to counteract any distraction from pushing the changes:

Date: Wed, 27 Mar 2024 12:46:32 +0000 From: krygorin4545 <krygorin4545@proton.me>

Subject: Re: Bug#1067708: new upstream versions as NMU vs. xz maintenance

Instead of having a policy debate over who is proper to do this upload, can this just be fixed? The named maintainer hasn't done an upload in 5 years. Fedora considered this a serious bug and fixed it weeks ago (<https://bugzilla.redhat.com/show_bug.cgi?id=2267598>). Fixing a valgrind break across many apps throughout Debian is the priority here.

What NeXZt?

Clearly social engineering techniques have much lower technical requirements to gain full access to development environments than what we saw with prior supply chain attacks like the Solarwinds, M.E.Doc ExPetya, and ASUS ShadowHammer incidents. We have presented and compared these particular supply chain attacks, their techniques, and their complexities, at prior SAS events [registration required], distilling an assessment into a manageable table.

Unfortunately, we expect more open-source project incidents like XZ Utils compromise to be exposed in the months to come. As a matter of fact, at the time of this writing, the Open Source Security Foundation (OSSF) has identified similar social engineering-driven incidents in other open-source projects, and claims that the XZ Utils social engineering effort is highly likely not an isolated incident.