

Table of Contents

The Blog	2
<i>Tony's Proxmox Reinstall and Migration Guide</i>	2
Use case	2
The solution	2
Traps	4
<i>Weird USB Bluetooth Music receivers</i>	5
<i>What is wrong with the 2023 League worlds opening ceremony?</i>	8
<i>Thinkpad T14 Gen2 AMD eDP Screen/Panel Upgrade</i>	10
Original Panel	10
New Panel	10
Worth it?	10
<i>CVE-2022-0847 - Dirty Pipe Vulnerability in Linux</i>	11
Introduction	11
Some background information	11
What are the causes?	12
Why is it important?	13
Who is affected and in what ways?	14
How can similar problems be prevented in the future?	14
If such a problem happened, how can we mitigate it?	15
Conclusion	16
Bibliography	16
Further reading	16

The Blog

Tony's Proxmox Reinstall and Migration Guide

Use case

There are many guides for doing “migrations” or “reinstalls” of proxmox online, sometimes with backups, sometimes with live copies etc.

My needs

I have an existing Proxmox 7 server running with a single boot SSD using LVM for PVE. I have additional SSDs for hosting my LXC and VM images, and hard drives that are bind-mounted into the LXCs where bulk storage is needed (eg: Immich and Nextcloud data directories).

While this has backups of the containers, (and the bulk storage is RAID Z1), I am a single SSD failure away from the system not booting, which I consider catastrophic when it is a 6 hour \$1000 flight away (thank Air Canada for that one).

I would like to reinstall PVE on this system, using two SSDs with ZFS RAID 1 for PVE, as 1) ZFS fits in far more with my current setup handling snapshots seamlessly, and 2) it gives me redundant boot and EFI partitions should one of the SSD's outright fail.

What is different

- I have NO quorum, I use a single instance of PVE.
- I can NOT do a live migration, as I have a single host, and will be reusing the hardware.
- I do NOT want to install all my LXCs and VMs from a backup - they live on a separate SSD I will not be wiping, and I would like to reuse those pools as is.
- I do NOT want to backup and restore PVE, I am changing SSD configurations (1 drive LVM to 2 drive ZFS)
- My old setup had many levels of jank during a PVE 7 to 8 upgrade, avoid re-using unnecessary old configs

Again, what I want, is to reinstall PVE onto a new SSD with a new filesystem, reuse my existing VM files and pools already on the drives. No backup/restore of LXCs. Just a simple, take everything, plunk it into new PVE.

The solution

Turns out the solution is actually relatively simple.

Since I tried a few things first and re-did this process twice, let me put some thoughts:

- Proxmox reads the configs for your VMs and LXC's from `/etc/pve/nodes`.
- I setup a full Proxmox Backup Server as an LXC another PVE instance (different location - not clustered)

The easy way

1. Make backups of everything you care about.
 1. IF YOU WILL BE RE-USING SSD'S, MAKE SURE YOU DD A BACKUP OF THEM!
 1. This one saved me and allowed me to repeat this procedure twice.
2. Grab the important config files you will need from the old host
 1. `/etc/pve/hosts`
 1. Ignore the keys and stuff, it'll cause pain and misery
 2. `/etc/wireguard/wg0.conf`
 1. If you rely on wireguard and want to be lazy and re-use old keys (not best security practice)
 3. `/etc/fstab`
 1. If you have non-standard mounts (I had an older btrfs HDD not in an zpool or lvm)
 4. `/etc/exports`
 1. If you have custom NFS exports to go from PVE over to VMs (LXC's just bind mount)
 5. `/etc/network/interfaces`
 1. If you have custom networking or VLAN setups
 6. `~/.ssh/authorized_keys`
 7. Any special configs you have.
3. Swap out the SSDs for the new ones, prepare proxmox installation media
4. Install Proxmox fresh on your two new SSDs!
 1. You should be able to boot into your new system
5. Bring your system roughly in line with your old setup. For me, this included
 1. Re-setup network interfaces and bridges the way they were
 2. Restore my fstab, wireguard, exports
 1. Wireguard needs apt install wireguard wireguard-tools
 2. `systemctl enable wg-quick@wg0.service -now` (start your tunnel at boot)
6. Re-import your storage devices (eg: my VM/LXC SSDs, data HDDs)
 1. Make sure they are all detected under disks!
 2. For my data HDD zpool, this was two steps
 1. `zpool import tonydata`
 1. causes it to show up in node → disks
 2. add it in the proxmox UI
 1. datacenter → storage → add new → zfs
 3. For my LVM SSDs, this was 1 step as they were already listed under disks
 1. add it in proxmox UI
 1. datacenter → storage → add new → lvm thin
 4. For my old BTRFS HDD, this was 3 steps
 1. Restore my old fstab (verifying UUID stayed the same)

2. mount -a
3. datacenter → storage → add new → btrfs
 1. Remember to choose the types of storage this is (backup and data, not disk images in my case)
5. Export my NFS pools (this one I forgot to do until my VM was mad)
 1. apt install nfs-kernel-server
 2. restore /etc/exports
 3. exportfs -av
7. Note, if you miss the above, it's probably not that bad, you'll be warned if things fail to start, it's iterative.
8. Restore /etc/pve/nodes/pve/lxc and /etc/pve/nodes/pve/qemu-server
 1. You should now see your left bar populate back up with all your past lxc and VMs
 2. Try starting them, maybe they work, maybe they error, if so look back to above.

That was it. Didn't end up using any of my backups from the backup server (surprisingly). Only used losetup on my dd to grab config files for networks and such retroactively to remember names and addresses.

Traps

- Do NOT, try to blindly restore your /etc/pve
 - Doing so will leave your /etc/pve in a bad state with keys that are lingering
 - You will NOT be able to log in and will need to systemctl stop, clear /etc/pve and such
- Connection error 401: permission denied - invalid PVE ticket
 - See my above point
- Do not try to RSYNC into /etc/pve
 - It will fail. It is a FUSE mount and will refuse the temp files rsync makes.
 - Use scp or an rsync flag to go direct to output file.
- Do NOT try to restore your PVE host from Proxmox Backup Server
 - Restoring the backup will yield errors with file overwrites
 - Allowing overwrites will break your configs in two and fail on /etc/pve
 - NOTE: proxmox-backup-client doesn't even backup /etc/pve by default (it only does root) so it's no use if you forgot that anyway
- PROXMOX BACKUP SERVER IS A RED HERRING!!
 - In fact, if you don't need to back up and restore VMs, and you've taken a DD of your old SSD, you don't need it.
 - it is excellent software for offsite backing up items, just not my current use case
- If it's not obvious, if you are using drive letters, like /dev/sda, sdb, YOU NEED TO STOP.
 - Please just use UUIDs, or heck even labels. After you reinstall, things WILL be shifting!!

In fact, let me do a whole segment on this

Proxmox backup server?

There seems to be some confusion with how this works, as a backup server can be added in two ways.

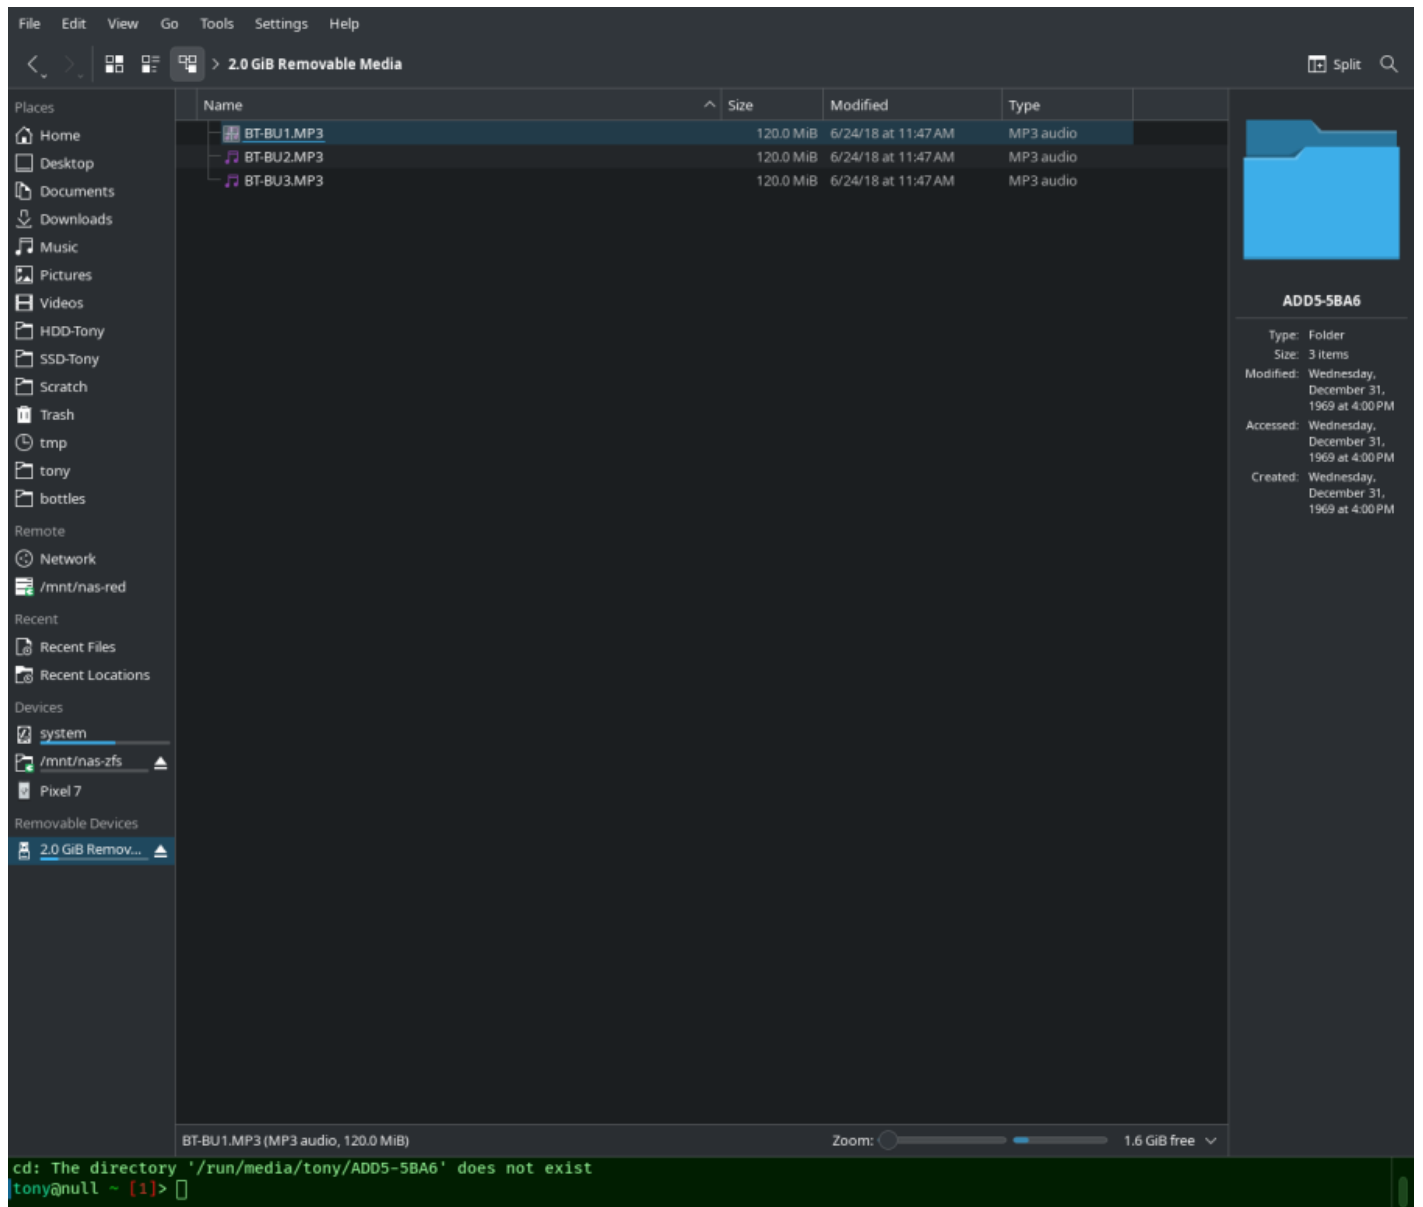
1. Datacenter → Storage
 1. This can be used for backing up the LXC and VMs on your system.
 1. Note: VM and LXC items backed up this way will keep some config data without /etc/pve, specifically ram and cpu amounts
 2. This can not be used to backup PVE itself
2. proxmox-backup-client
 1. This can be used to backup PVE. by default, backs up a partition of your choosing
 2. If you backup /, make sure you also include /etc/pve which is a fuse mount
 1. If you already DDED the drive, you can always losetup and mount the disk image instead of using pxar.

2025-01-03 05:22 · Tony

Weird USB Bluetooth Music receivers



My USB powered bluetooth audio receiver STARTED SHOWING UP AS A MASS STORAGE device.
I DIDN'T EVEN KNOW THAT THING HAD DATA LINES?!?
Did I get rubber-duckied by a BLUETOOTH RECEIVER?
WAIT YOU CAN MOUNT THE STORAGE DEVICE?!?



The screenshot shows a file manager window titled "2.0 GiB Removable Media". The main pane displays a table of files:

Name	Size	Modified	Type
BT-BU1.MP3	120.0 MiB	6/24/18 at 11:47 AM	MP3 audio
BT-BU2.MP3	120.0 MiB	6/24/18 at 11:47 AM	MP3 audio
BT-BU3.MP3	120.0 MiB	6/24/18 at 11:47 AM	MP3 audio

The right sidebar shows details for the folder "ADD5-5BA6":

- Type: Folder
- Size: 3 items
- Modified: Wednesday, December 31, 1969 at 4:00 PM
- Accessed: Wednesday, December 31, 1969 at 4:00 PM
- Created: Wednesday, December 31, 1969 at 4:00 PM

The terminal at the bottom shows the following command and error:

```
cd: The directory '/run/media/tony/ADD5-5BA6' does not exist
tony@null ~ [1]>
```

WTF IF YOU PLAY IT, THE FILE HAS LEGIT MP3 TAGS

```

File Edit View Bookmarks Plugins Settings Help
tony@null ~> ls -lah /run/media/tony/ADD5-5BA6/
total 361M
drwxr-xr-x  2 tony tony  32K Dec 31  1969 ./
drwxr-x---+ 3 root root   60 Jul  7 15:02 ../
-r--r--r--  1 tony tony 120M Jun 24  2018 BT-BU1.MP3
-r--r--r--  1 tony tony 120M Jun 24  2018 BT-BU2.MP3
-r--r--r--  1 tony tony 120M Jun 24  2018 BT-BU3.MP3
tony@null ~> mpv /run/media/tony/ADD5-5BA6/BT-BU1.MP3
[ffmpeg/demuxer] mp3: Estimating duration from bitrate, this may be inaccurate
(+) Audio --aid=1 (mp2 2ch 48000Hz)
File tags:
Artist: btdongle
Album: bt dongle
Date: 2019
Title: BT MUSIC1
AO: [pipewire] 48000Hz stereo 2ch s16p
A: 00:00:00 / 01:14:53 (0%)
Exiting... (End of file)
tony@null ~> mpv /run/media/tony/ADD5-5BA6/BT-BU1.MP3

```

```

File Edit View Bookmarks Plugins Settings Help
libavcodec      61.  3.100 / 61.  3.100
libavformat     61.  1.100 / 61.  1.100
libavdevice     61.  1.100 / 61.  1.100
libavfilter     10.  1.100 / 10.  1.100
libswscale      8.   1.100 / 8.   1.100
libswresample   5.   1.100 / 5.   1.100
libpostproc     58.  1.100 / 58.  1.100
[mp3 @ 0x77b664000c80] Estimating duration from bitrate, this may be inaccurate
Input #0, mp3, from '/run/media/tony/ADD5-5BA6/BT-BU1.MP3':
  Metadata:
    title       : BT MUSIC1
    album       : bt dongle
    artist      : btdongle
    date        : 2019
  Duration: 01:14:53.88, start: 0.000000, bitrate: 224 kb/s
  Stream #0:0: Audio: mp3 (mp3float), 48000 Hz, stereo, fltp, 224 kb/s
^C
tony@null ~ [123]>

```

OH SHIT I Think this is for it to work on crappy stereos that just play MP3 from a flash drive!!!
 To adapt them to work with bluetooth
 hm I don't hear anything and the mp3 file exits though
 so I guess you need a dumber player that doesn't try to seek in a file
 WAIT I THINK I KNOW WHAT THE OTHER FILES ARE
 YOU CAN'T PLAY THEM, BUT IF YOU OPEN IT, IT TELLS THE PHONE TO GO TO THE NEXT SONG
 so it's used to detect when you pressed previous/next!!! and from that, it signals to the bluetooth device

to go back/forward

who the f*** makes this ASIC?!? There's no way this thing is economies of scale

2024-07-07 22:18 · Tony

What is wrong with the 2023 League worlds opening ceremony?

The original league worlds opening ceremony video from this year has very many technical quirks. Here is what I noticed and spent a few (prolly a dozen) hours fixing.

- The audio on the YouTube upload is inexplicably in mono and sounds notably worse than the Twitch livestream which was in stereo.
- Likely destructive interference when collapsing stereo to mono?
- The whole 'scripted' section of the opening is actually running at 29.97 fps, and interpolated to 60 but in the worst way.
- Half of their cameras are running at 29.97 which means every frame is duplicated for the stream which is 59.94 - this on it's own is fine. Most of the side-angle cameras are 29.97.
- The other half of the cameras are also running at 29.97, but are interpolated to 59.94 THROUGH FRAME BLENDING?!?
- This means that every other frame is literally a 50/50 blend of the two frames before it. This is what causes that motion-blur like feel, and you can instantly tell which cameras they are on.
- It might be the cameras that also do AR are doing blending? unsure.
- In fact, the main stage camera and the cable cam blend on opposite frames - eg: one has good even frames, blended off, the other is vice versa. This might be because of an internal delay? I'm not sure. Don't get me wrong 1080p29.97 is fine, although odd since sports is usually 720p60 or 1080i60.
- The version on YouTube took a video that had TV black levels, mapped it to PC, then was converted to TV again, resulting in crushed black and white levels losing detail
- The version on Twitch has the audio and video about 100 ms out of sync, this is most noticeable with the fireworks at the end of gods not being in line with music. Interestingly, the audio and video ARE synced correctly on YouTube version.
- TV (and most video formats) uses 16-235 to express black-white instead of 0-255, whereas PC is 0-255. So, if you take a 16-235 feed, map it to 0-255, then crop it into 16-235 again, you effectively increase contrast and reduce your dynamic range, that is what seems to have happened on YouTube.
- Another odd quirk: some of the camera angles are different between the YT and Twitch versions? Either these were cut after the fact, or it was filmed weird? Not sure, I can upload a side by side comparison if anyone is curious. Notable on some audience and a few side angle shots.
- Also if you watch the side by side, depending on if you sync up the video or audio feeds, it almost looks like there are 2 video systems, and depending on which you sync up on, some transitions happen 100 ms before or after the other, with the other half being synced. I think the AR stuff is probably separate?
- About that frame blending, there are 4k60 audience camera recordings on YouTube, and if you watch those carefully, you'll notice the camera feed going to the video walls also does frame blending. This tells me that whatever system video was going into first was doing the blended

frames, but again, only on half of the cameras.

So, to fix this

- I took the stream from Twitch, split the audio and video. For the video, I couldn't simply drop all odd or even frames because their different cameras interpolate different sets of frames.
- From my cut, the "main" stage camera had odd frames correct, and even frames blended, and the cable cam that is further away had all even frames correct and odd frames blended.
- So I spent 3 hours clipping each camera angle, and dropping either the odd or even frames accordingly.
- I then re-synced the audio with that 100ms offset. This left me with a 29.97 fps video. That is my "cleaned up" version that I might upload, but, since the original was 60 fps, I wanted to interpolate it, but properly, not through frame blending.
- So I ran the whole thing through an AI video interpolation tool overnight on my GPU, and there you have it, 4k60 (technically 59.94). I'll add the 29.97 maybe as well.

Stuff I did not fix:


- The yellow jacket intro video is filmed in 24 fps, and brought up to 30, but not in a very consistent way, so it's staying as is.
- I can't change mic volumes, there isn't a lot, but you can pretty easily tell which parts are lip-synced anyway (listen for the reverb different, or when vocals sounds like they are stereo - eg: first 10s of vocals in gods sounds live, most of the rest is off the tape).
- I also can't fix their broken shadows on their AR characters.

Note: I do not own the music or video, those are from Riot Games. I'm simply uploading this in case anyone else wanted to watch it without the bizarre quirks of the version on YouTube.

If anyone from the Riot Games technical team is watching, I'm genuinely curious about how their production setup works and I would love to chat. (Or if you know someone who works with the technical team)

This is not meant to be a dunk on the performance, overall it was well put together, these are just 'quirks' I saw and couldn't unsee.

Here is the fixed version



2023-12-20 07:36 · Tony

Thinkpad T14 Gen2 AMD eDP Screen/Panel Upgrade

If you read [Thinkpad T14 Gen2 AMD](#), I hated the screen they shipped with the T14, and thought it was a disgrace for a “premium” laptop.

Well, it annoyed me so much that I started to research how to upgrade the screen of the laptop.

Original Panel

The original panel is a was an Innolux N14HCA-EAE.
It is IPS-type, <250 nits, and has garbage colours.

Since it is 1080p non-touch, replacing it with another 1080p panel isn't too hard.
If you had a 4k screen, you would need a new eDP cable. If touch, replacing the digitizer is another step.

Look for the Thinkpad maintenance guide for removing an old screen. Follow it like Lego.
This is an irreversible process. Save the bezel. It's supposed to be replaced, good luck finding a replacement. Pry it slowly.

The original screen is held by double sided tape. I didn't replace mine, you maybe should - again, good luck.

New Panel

I bought the Innolux N140HCG-GQ2.

It's low power, 400 nit, 90+% sRGB, IPS-like, 1080. Much better.

No new cable needed, drop in replacement basically. Comparison sites label it the best of the 3/4 “low power” display options.

You can find via AliExpress and hope it's not a fake.

Re-use existing tape and bezel when reassembling, clip display properly.

Worth it?

100%. The brightness and colour gamut are very much noticeable and appreciated.
It makes me actually enjoy looking at the laptop screen.

If you have the patience and courage, go for it! LMK how it went!

2023-07-02 15:42 · Tony

CVE-2022-0847 - Dirty Pipe Vulnerability in Linux

Introduction

The Linux Dirty Pipe vulnerability, also known as CVE-2022-0847 is major a vulnerability first discovered near the end of February 2022 which affects Linux kernel versions 5.8 and above. The Dirty Pipe vulnerability is considered severe since there aren't workarounds (besides upgrading to a patched version of the Linux kernel) and it allows local privilege escalation allowing users to write to files they don't have write access to (such as `/etc/passwd`), which then allows them to run commands as root or drop to a root shell.

Before I go any further, if you are reading this post from a computer running Linux (or GNU/Linux), make sure you have updated your kernel to 5.16.11, 5.26.36 or 5.10.102 to make sure you're protected against this vulnerability!

This vulnerability was first discovered last month when investigating an older bug report regarding corrupted files. In particular, there were reports of corrupted CRC checksums at the end of zip files in a particular use case involving splices and pipes.

Some background information

There are a few background elements that I think are useful to understand this vulnerability and how it works: Pipes, pages, the page cache and the splice system call on Linux.una20

First: Pipes. Pipes in Linux are a unidirectional, inter-process communication method provided by the kernel. Pipes are commonly used to pass information between 2 applications. The output of the first application is saved to the pipe, which is then read by the second application. Pipes can be used both in the shell using `|` to redirect the output of one program to the input of the next (such as if you do `cat foo | grep a | cut -c 2- | uniq | sort | tee bar`), or within programs, which is more flexible and powerful. Under the hood, pipes are treated and behave just like files, which is part of the Unix philosophy. However, unlike files on secondary storage, the data in pipes is stored in a buffer (that works like a ring using pages) instead of on the disk, where data from the first application is added to the buffer, and the second application reads from it. In other words, it's a handy way to send data from one application to another, where one pushes data into the pipe, and the other pulls it from the pipe.

Next: Pages. As discussed in the earlier modules of the course, pages are blocks of virtual memory. Pages are generally the smallest unit of data used for virtual memory. These pages are what will contain files, data, code, so on, and can be loaded from the secondary disk (SSD/HDD), or from system memory. On Linux based systems, the page size is typically 4 KiB (4096 bytes). (You can check the page size for yourself by running `'getconf PAGESIZE'`)

To speed up I/O operations when a user is frequently working with certain pages, Linux maintains a page cache, which will essentially cache the commonly used items from secondary memory (SSD/HDD) onto primary RAM. This means that if you're continually loading a part of a file, instead of waiting for a hard drive each time, it can be cached and retrieved from the cache in memory. This cache is usually in

unused parts of memory, and can be easily cleared if an application needs to actually use the memory. Pages in the page cache that are modified while in memory (such as by a write to the file) will mark the page as 'dirty', which have to be written back to the disk to save/persist those changes.

Finally, the `splice()` system call. `Splice` is a system call that moved data between files on the disk and pipes without having to load and go through the userspace. Generally, `splice` can work by just remapping pages, which makes it fast. This means that instead of passing the page around, it passes page references, without having to do the copy. You can `splice` a file to a pipe, then `splice` back from the pipe into another file without dropping to userspace, which makes this both fast and secure.

What are the causes?

The dirty pipe vulnerability takes advantage of a pipe and a `splice` call. This started as a bug that was introduced with a commit which refactored the pipe buffer code, changing how 'mergeable' checks were done.

As mentioned in the section above, all data, including the data temporarily stored in a pipe lives in a page, which can get cached into memory. The first write into a pipe allocates a 4 KiB page. Then, this page is used for writes until it fills up, at which point a new page is allocated. However, if you use the `splice` system call, the kernel first loads the data from the file (source) into the page cache, then, creates a 'pipe_buffer' struct (the destination pipe) which points to the page cache. This is a zero-copy, since it essentially just copies the reference to the page in the page cache. However, even if the page isn't full, it can't be appended to, since the page is owned by the page cache, not the pipe.

The checks to ensure that data isn't appended to page in the page cache have changed over time, but the latest revision with Linux 5.8 turned a bug into a critical issue. Since kernel 4.9, there has been a bug causing the flags member of the pipe_buffer struct to be missing, however, the new critical issue allows a user to inject `PIPE_BUF_FLAG_CAN_MERGE` into the page cache reference, which then allows the user to overwrite data in the page cache! Once the `PIPE_BUF_FLAG_CAN_MERGE` flag is injected into the reference, writing new specially crafted data into the pipe that was originally loaded with the `splice` causes the kernel to append this data to the end of the page in the page cache, instead of creating a new page.

Since the `splice` operation loaded the contents of a file into the pipe, by placing the pages for the file in the page cache, if we can write to the page cache, we can effectively write to the version of the file that will be used by the OS. This is critical, since you only need read-access in order to load the contents from a file into a pipe. However, as seen here, by setting the flag, we can write to the page in the page cache. This means that any user can write into files they would normally only have read access to. Essentially, the kernel is tricked into thinking this is a normal page for the pipe that is safe to write to, however, it's actually the contents of other files. Yikes.

Now, there is one aspect of this to consider. By using this vulnerability, you can write data into the page cache, which is loaded with the contents of another file. However, since the kernel is unaware of this change, the page in the cache is not marked as dirty. This means that eventually, either on system restart, or if the kernel needs to reclaim the memory, the page in the cache will be dropped. This is why this bug was hard to spot in the wild, since in most cases, it would be gone by the time the system restarted. On the other hand, since the data in the page cache isn't written to the disk, and lasts for a

considerable amount of time, this provides a stealthy attack surface that doesn't leave a trace on the disk.

Why is it important?

This vulnerability is critically important, since it can be exploited to achieve local privilege escalation. This vulnerability allows any unprivileged user to write to pages in the page-cache which correspond to files they only have read access to. Since the kernel will use the page cache to load a file if it is cached, this effectively allows the user to overwrite data into any file they have read access to, with some limitations.

This can be exploited to get root access (or a root shell) through several means. For example, through a few steps, one can simply clear the root password from `/etc/passwd`, add SSH keys, or change data used by root processes. The simplest is likely by disabling the password for the root user. While normal users can't access `/etc/shadow`, by removing the `x` in `/etc/passwd`, you can remove the lookup to `/etc/shadow`, and allow password-less login to the root user. For those who are curious, there is a proof-of-concept exploit available on the original dirty pipe vulnerability page that allows writes into read-only files. Once an attacker has root access on your system, there's little they can't do.

A system with this vulnerability can be exploited by attackers for any malicious purpose following a few steps. First, the user creates a pipe. This is a pipe which they must have write access to. Next, you fill the pipe with random or arbitrary data. The contents don't matter, it just needs to fill the pipe so that the `PIPE_BUF_FLAG_CAN_MERGE` which is later incorrectly checked will get set in all of the pages and entries used by the pipe. Next, you empty the pipe, by reading the data from it. This leaves the flag set in all instances of the `pipe_buffer` struct in the ring.

For step 4, you need to make a `splice` system call for the file you want to modify, opened with `O_RDONLY` (read only), and splice it into the pipe just up to the target offset. The target offset is the location in the file where you want to write your new data to and make your changes. For example, this could be the location of the root user password. Next, you write the data you want to replace in the file into the pipe. As discussed in the last step, instead of creating a new anonymous `pipe_buffer` struct with new pages, the new data added to the pipe will begin to replace the data in the page from the read-only file (which is in the page cache) because the `PIPE_BUF_FLAG_CAN_MERGE` flag is set. Recall that that flag would normally only be set if this were a normal pipe, but remained set.

This exploit was particularly interesting (and terrifying) for myself since I had my secondary computer running an earlier build of Linux 5.16, which was vulnerable, and I was able to test both the small code by the original author, as well as their proof-of-concept exploit which both worked. While not trivial, it is a vulnerability that is relatively simple to exploit. While there was somewhat similar exploit named Dirty Cow from a few years ago, this is simpler to exploit.

There are however a few limitations to these exploits, although they are all simple to overcome once you gain root access or a root shell. First, you need to have read access to a file. Without read access, you won't be able to use the `splice` call. Next, the offset that you want to overwrite can't be on a page boundary (ie: can't be a multiple of 4096). This is because you need to splice at least 1 byte from the target file into the pipe. Next, you are limited to overwriting 1 page (4096 bytes). While this might seem short, since you can't change large files, it's more than enough to update important files such as

/etc/passwd. Finally, you can't resize the file, so the data you write can't alter the file size. Again, once you have root access, the other limitations wouldn't matter for an attacker, since they can use their root access to change any other files they wanted.

Who is affected and in what ways?

Any device running Linux 5.8 or later (without the latest patches) are affected. This includes many servers, desktop computers and laptops, as well as many recent mobile devices running Android. In particular, systems running many major distributions from the past 2 years without the most recent updates are vulnerable. This includes Debian 11 (bullseye), Fedora 33 and up, Ubuntu 20.04 and up, as well as many others including Arch, Gentoo, OpenSUSE, as well as their derivatives. This constitutes a good portion of server and desktop systems running recent installations of GNU/Linux.

Although it is a local vulnerability, meaning that it doesn't add a new vector for an outside attacker, if an attacker has any access to a user on the system, be it physical or digital such as over SSH, Telnet or VNC, they can exploit this vulnerability to gain root access.

Equally important is the use of the Linux kernel in the Android mobile operating system. Within a few days of the discovery of the original vulnerability, the bug was reproduced on a Google Pixel 6, and reported to the Android Security Team. If the bug is reproduced, then it means that there will also be exploits targeting these devices. Of course, this impacts more devices than just the Pixel 6. For example, the Samsung Galaxy S22 series are all based on Android 12, running Linux kernel 5.10.43, which is vulnerable. Android versions prior to Android 11 used Linux kernel versions 4.4, 4.9, 4.14, 4.19 or 5.4, however, starting with Android 12, support for Linux 5.10 was added, with Android T (the experimental branch) also using 5.10. This means that many flagship phones are going to ship with a build of Linux 5.10. Given the poor record of Android manufacturers issuing software updates, this can end up affecting a good number of recent phones that were released.

How can similar problems be prevented in the future?

Given that this bug was introduced due to programmer error when refactoring a section of the code, it's harder to provide many meaningful suggestions. A solution that is frequently presented is to have a more stringent code-review process, with more developers or a stronger emphasis on security. However, patches going into the Linux kernel already undergo a thorough code-review from multiple developers. While there are also suggestions to avoid refactoring some code since it can introduce bugs such as this one, refactoring is also important in order to clean up the code and keep it consistent.

Some mistakes like this could be caught in part by using analysis tools. For example, the basis of this vulnerability was originally a bug with uninitialized flags for the struct. This type of mistake such as uninitialized variables in structs could be caught if the team added corresponding checks. However, this of course comes with the downside that the developers might also get bogged down with tons of false positive reports. Some have also suggested switching to a safer language such as Rust, but in cases like this where it's very specific behavior in specific calls and specific checks, which don't necessarily break the safety of the language, it's debatable it would have caught a mistake such as this one.

We have also seen that having a simpler codebase makes auditing much simpler while also making it

both faster and safer since it becomes easier to stop mistakes. For example, the Wireguard VPN protocol which was written to be a faster, safer VPN protocol was welcomed into the Linux kernel in that past years since it's so much simpler then competing protocols such as OpenVPN (which is based on the comparatively massive OpenSSL and TLS libraries). However, a kernel is much more sophisticated than VPN protocols, and given that Linux is a monolithic kernel, it is much harder to just have simpler code - although there are projects such as Redox OS which are both supposed to be simpler, using a microkernel architecture, as well as being written in the safer Rust language.

Overall, this is one of the aspects where open source shines. We can always have more developers looking at and through the code, and it allows users such as Max Kellermann (the original author of the vulnerability) to not only find bugs such as this one, but also patch them in a timely manner.

If such a problem happened, how can we mitigate it?

This vulnerability was reported professionally by the security researcher, who alerted the Linux kernel security team within a day of his discovery of the vulnerability, along with a patch. This was quickly sent to the Linux Kernel Mailing List the next day, where it was merged into the next release within a few days. Overall this problem fixed very quickly by the Linux developers, with patched versions of the kernel out within a week of the discovery of the original vulnerability. This is excellent, and it shows that security issues such as this one are taken very seriously by the developers who can work through these issues and patch them in an effective manner.

Once again, this is a good example of how the open source software community is very good at responding to issues like this once they are found with everyone working together to get it patched and ensure systems get updated. This vulnerability was fixed for the release of versions 5.16.11, 5.15.25, and 5.10.102. Of these, 5.16.11 is the latest stable release that distributions like Fedora and Arch use, with 5.15 and 5.10 being LTS releases.

Once the new versions of the kernel were released, all of the distributions mentioned in my section above have either upgraded their kernel to include the patch for this vulnerability, or, backported the patch for this vulnerability into their kernel versions. This is a good example of the author, kernel developer, as well as distribution maintainers working together to ensure security issues like this get patched.

As I mentioned in the start of this post, the mitigation for this vulnerability for both desktop and server systems is to update your kernel! Whether you're maintaining a server or your system, make sure your software is up-to-date! For servers in particular, make sure you're running a version of your distribution that is actively maintained, such as an LTS release to ensure patches like this one get backported to ensure your systems are safe.

Mobile users on the other hand will have to wait for updated versions of Android to be released for their devices, likely through security updates. We can hope that they will be as responsive as distribution maintainers are, and apply these patches and deliver new updates in a timely manner. Since this bug was introduced in Kernel 5.8 from a few years ago, and the most recent Android version to ship with a vulnerable kernel is Android 11, most devices that are vulnerable are relatively new, increasing the chance they will get patched.

Conclusion

I hope you enjoyed reading this and learned about the dirty pipe vulnerability on Linux, and how pipes, pages, the page cache and splice call work in general!

Bibliography

- “CVE-2022-0847,” CVE, 03-Mar-2022. [Online]. Available: <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2022-0847>. [Accessed: 14-Mar-2022].
- D. Das, “What is the dirty pipe exploit in linux and how can you fix it?,” Make Use Of, 10-Mar-2022. [Online]. Available: <https://www.makeuseof.com/what-is-the-dirty-pipe-exploit-in-linux-and-fix-it/>. [Accessed: 14-Mar-2022].
- D. LeClair, “New dirty pipe linux vulnerability is the worst in years,” How-To Geek, 08-Mar-2022. [Online]. Available: <https://www.howtogeek.com/790435/new-dirty-pipe-linux-vulnerability-is-the-worst-in-years/>. [Accessed: 14-Mar-2022].
- M. Kellermann, “The dirty pipe vulnerability,” The Dirty Pipe Vulnerability - The Dirty Pipe Vulnerability documentation, 07-Mar-2022. [Online]. Available: <https://dirtypipe.cm4all.com/>. [Accessed: 14-Mar-2022].
- P. Arntz, “Linux ‘dirty pipe’ vulnerability gives unprivileged users root access,” Malwarebytes Labs, 11-Mar-2022. [Online]. Available: <https://blog.malwarebytes.com/exploits-and-vulnerabilities/2022/03/linux-dirty-pipe-vulnerability-gives-unprivileged-users-root-access/>. [Accessed: 14-Mar-2022].
- S. Hazarika, “Linux kernel bug dubbed 'dirty pipe' can lead to root access, affects Android devices as well,” XDA Developers, 09-Mar-2022. [Online]. Available: <https://www.xda-developers.com/dirty-pipe-linux-vulnerability-root-android/>. [Accessed: 14-Mar-2022].

Further reading

The original post disclosing this vulnerability has lots of other technical information, including how the vulnerability was discovered due to a problem with repeated corruption on web-server logs.

<https://dirtypipe.cm4all.com/>

The commit that introduced this vulnerability is here:

<https://github.com/torvalds/linux/commit/f6dd975583bd8ce088400648fd9819e4691c8958>

For those of you that want to try the proof-of-concept exploit on your unpatched devices, the source code is available here: <https://packetstormsecurity.com/files/166229/Dirty-Pipe-Linux-Privilege-Escalation.html>

If you have any questions or comments, feel free to leave it below!

2022-03-14 07:06 · Tony

[Older entries >>](#)

From:
<https://wiki.tonytascioglu.com/> - **Tony Tascioglu Wiki**

Permanent link:
<https://wiki.tonytascioglu.com/blog>

Last update: **2023-12-20 07:34**

