

Additional Linux Commands & Guides

Additional troubleshooting commands & help guides for software installed on systems.

- [Small Commands, Issues And Guides - A Collection](#)
- [Transferring server data between two dedicated systems - Wget & Python3](#)
- [Cannot Access UDP Connected Games Due To OVH Firewall](#)
- [Accessing A Website Or Control Panel Shows "Too Many Redirects" Error](#)
- [Error 1014 Tebex - CNAME Cross-User Banned](#)
- [Why Is X Not Installed On The System - How Docker Images Work](#)
- [PHP Extensions Locations For .so Files](#)
- [Installing Modules On cPanel For WHMCS - IonCube, FileInfo & Soap](#)
- [LiteSpeed WebServer With Pterodactyl \(CyberPanel\)](#)
- [Searching Through Files - Finding A Particular String And Updating It](#)
- [Installing SQLite3 To SQL with Python3 Package To Convert SQL-Lite DB to MySQL for import](#)
- [ItemsAdder With CloudFlare Proxied Domain For IP Protection](#)
- [Get the “/dev/disk/by-id” information for FSTab from the udevadm command](#)
- [Resizing single volume group and logical volume after disk expansion](#)

- [Resizing a logical group then volume within a parent RAID number](#)
- [CloudFlare does not minify javascript correctly when working with Grafana](#)
- [Install Custom Version Of MySQL Server](#)
- [Networking With Hetzner - With and Without additional MAC's configured](#)
- [Running Lancache & Another Webserver On Top](#)
- [Installing Debian/Ubuntu & Then Disk Unmounts - Debootstrap Error](#)

Small Commands, Issues And Guides - A Collection

This page is for commands, issues or guides which don't need to be added as a separate page. You can search for them with **Ctrl + F** or click the "Search" at the top.

Only Allow PublicKey Authentication Overriding Password Auth

No idea why, some systems are completely ignorant and ignoring the demand for no password auth.

Add this line to bottom of SSHD config. No idea why, but the The PA parameter is no and ChallengeResponseAuthentication is also no, but still auths with password... **sigh.**

AuthenticationMethods publickey

Give A MongoDB User Permission To Command "mongodump"

use admin

```
db.grantRolesToUser( "USER", [ "readWrite" , { role: "readWrite", db: "config" } ] )
```

Remove Apache2/Nginx Python3's modules for Certbot

Apache2: python3-certbot-apache

Nginx: python3-certbot-nginx

Remove a package that is being stubborn and will not remove

```
sudo dpkg --force-all -P <package-name-to-remove>
```

Absolute headache with MySQL 8.0 mixing with MariaDB's installs

MySQL 8.0 has the config file in */etc/mysql/mysql.conf.d/mysqld.cnf* like MariaDB has */etc/mysql/mariadb.conf.d/50-server.cnf*.

BY DEFAULT MySQL binds publicly... MariaDB binds to internal localhost.

Use **mysqld --verbose --help | grep bind** to actually see it bind publicly on first installation...

If you mix MariaDB's config files + MySQL's, it's not my.cnf, mysql.cnf, the conf.d folder entirely or any other configuration. **It is literally mysqld.cnf in the mysql.conf.d folder.** What utter BS.

Fixing the repository error for changing it's "version" for updates

Example: *N: Repository 'http://deb.debian.org/debian bullseye InRelease' changed its 'Version' value from '11.7' to '11.8'*

One command: `apt-get --allow-releaseinfo-change update`

Taken from:

https://www.reddit.com/r/debian/comments/ca3se6/for_people_who_gets_this_error_in_release_changed/


The certbot command that actually works to change the email

`certbot update_account -email <email>`

SlimeWorldManager Installation

- Download the latest Spigot/Paper jar you require and pop it into the main directory.
- Put the slimeworldmanager-plugin-<version>.jar into the same place as the server.jar.

- Update your startup command to include **-javaagent:<thejar>**, like the



below:

Taken from: <https://github.com/cijaaimee/Slime-World-Manager/blob/master/.docs/usage/install.md>

Crons not running for WHMCS on cPanel - They hang and sit there doing nothing

- Log into WHM (The admin side of cPanel)
- Navigate to "MultiPHP INI Editor" by searching it on the left hand side.
- Select the version of PHP that your site uses, using the editor is easier.
- Increase the following variables:
 - "max_execution_time" - 240/300 should be fine.
 - "max_input_time" - same as above.
 - "memory_limit" - Higher than 1024M.
- Click "Apply"
- Repeat for all versions of PHP that are having a timeout error.
- Restart cPanel's FPM with systemctl restart ea-php81/82-php-fpm.

Taken from: <https://support.cpanel.net/hc/en-us/articles/360052237994-How-to-increase-the-Max-Execution-Time-directive-for-PHP>

RAID Ubuntu 20.04

Taken from: <https://kifarunix.com/setup-software-raid-on-ubuntu-20-04/>

Note - If you get error: *partition length of [BIGGER SIZE] sectors exceeds the msdos-partition-table-imposed maximum of [SMALLER SIZE]*

Follow: <https://askubuntu.com/questions/84538/trouble-creating-3tb-ext4-partition-due-to-msdos-partition-table-imposed-error>

Switching RAID 1 to RAID 0

Scroll down a bit to locate the answer from nmr, which includes removing the mirror, changing, waiting then resizing.

Taken From: <https://serverfault.com/questions/915284/is-it-possible-to-convert-raid1-to-raid0-without-system-reinstalation>

Remove .html/.php extension

Read the answer from Arnon, the big threaded part with explanations.

Taken from: <https://stackoverflow.com/questions/38228393/nginx-remove-html-extension>

Can't run a module or build a NodeJS app due to no such file and has no installation candidate

```
Error: libnode.so.72: cannot open shared object file: No such file or directory
  at Object.Module._extensions..node (node:internal/modules/cjs/loader:1280:18)
  at Module.load (node:internal/modules/cjs/loader:1074:32)
  at Function.Module._load (node:internal/modules/cjs/loader:909:12)
  at Module.require (node:internal/modules/cjs/loader:1098:19)
  at require (node:internal/modules/cjs/helpers:108:18)
  at Object.<anonymous> (/var/www/minetrack/node_modules/sqlite3/lib/sqlite3-binding.js:4:17)
  at Module._compile (node:internal/modules/cjs/loader:1196:14)
  at Object.Module._extensions..js (node:internal/modules/cjs/loader:1250:10)
  at Module.load (node:internal/modules/cjs/loader:1074:32)
  at Function.Module._load (node:internal/modules/cjs/loader:909:12) {
  code: 'ERR_DLOPEN_FAILED'
}
```

```
root@ns1009801:/var/www/minetrack# sudo apt-get install libnode72
Reading package lists... Done
Building dependency tree
Reading state information... Done
Package libnode72 is not available, but is referred to by another package.
This may mean that the package is missing, has been obsoleted, or
is only available from another source

E: Package 'libnode72' has no installation candidate
```

One simple command to clear out the old NPM packages already downloaded from cache.

npm clean-install

You can then go about reinstalling with yarn build or another method used.

Any NPM package doesn't install into /usr/local/bin, such as ghostcms - Updating the prefix location due to weird bug

Taken from: <https://askubuntu.com/questions/1102579/using-npm-to-install-file-to-usr-local-bin>

For some weird reason, on some installs of NodeJS and then installing npm, it ends up at /usr and not /usr/local to allow /usr/local/bin packages to be viewable, such as GhostCMS's ghost-cli. Stupid bug, no idea why. Run the set prefix command to fix it, then install Ghost's CLI fine.

The reason why this needs doing is due to the error "Error: Cannot find module '/usr/local/bin/ghost'" and "MODULE_NOT_FOUND" which is absolutely idiotic since you literally just installed it, but in the wrong place. This will happen when trying to start the service and it fails.

```
root@S01: /var/www/website.com# npm get prefix
/usr
root@S01: /var/www/website.com# npm config set prefix /usr/local
root@S01: /var/www/website.com# npm get prefix
/usr/local
sudo npm install ghost-cli@latest -g
```

Date command in linux - Echo and crontab support

When using the date command in terminal, if you are echoing the line, you need to make sure you don't escape the percentages as it will include the slashes.



- echo \$(date +"%d-%m-%Y").txt




- echo \$(date +"%d-%m-%Y").txt



However, when inside crontab, you will need to escape them otherwise cron believes it is an operand and not a variable to pass. An example is below where you also need the path in a full string for it to echo properly and understand the date variable.

-  image not found or type unknown
-  image not found or type unknown

You can also use **-d** for naming how we name things, such as "yesterday" "tomorrow" or "2 months ago".



- **Note:** You can also do broken English such as "2 month ago" instead of months!
- **Taken from:** <https://stackoverflow.com/questions/17015187/how-does-date-d-parse-dates>

Transferring server data between two dedicated systems - Wget & Python3

The easiest way to transfer a tarball or zip file from one dedicated system to another is by using wget and a python3 webserver, protected from sending out the data from UFW allow in only that IP address or using the internal network if your system has it, such as OVH's vSwitch/vRack system.

Open up SSH. In this example, we will be using Pterodactyl's UUID's in the volumes folder, but this also applies to any zip/tarball you want to do this on such as servers stored in /home or /srv.

Go to `cd /var/lib/pterodactyl/volumes` and then do `ls` to see the containers.

```
root@S01:/var/lib/pterodactyl/volumes# ls
031d5d21-2d7a-4381-81ae-3b4602a14985  70409e14-a41c-4c37-910c-cfe364ed01d5
2df033ff-c2c4-4f62-b421-bace8040a03c  8d974c46-a6b7-4297-b103-c360603b7cfd
3af93aa1-8752-4453-befe-8171fdcaa63a  93d0d9a7-ff59-4714-8430-1d1729d22c0c
3e21c8fa-c0f3-4208-ba3b-b20f07da7d89  ba140f0c-50d0-422c-a185-48d86bded8f8
608b0362-f2e5-46c3-949f-64f8b420afed  bb99bfb3-3357-485a-ad1f-5a5be7e18e6e
6a348c25-eda1-4896-9892-86d10099ab45  feadc8c2-ec1e-42ff-bf3b-0ce951f36c4f
root@S01:/var/lib/pterodactyl/volumes#
```

We will use "031d5..." as an example. Run `tar -zcvf SERVERNAME.tgz 031d5` then press tab to autofill the entire server ID.

```
# tar -zcvf hub.tgz 031d5d21-2d7a-4381-81ae-3b4602a14985/
```

Let it finish tarballing. You'll know when it's finished as the output stops showing you what files it's archived and goes back to the normal terminal line waiting input.

```
031d5d21-2d7a-4381-81ae-3b4602a14985/help.yml  
root@S01:/var/lib/pterodactyl/volumes#
```


Do `ip a` on your destination system and find the public IP. Remember, some systems have internal networking so you might be able to do `wget 10.10.10.1:8000/hub.tgz` for example instead of the public IP and whitelisting it, as it normally is already whitelisted.

Image not found or type unknown




Once you've found either the public or the private IP, go back onto your source system and type `ufw allow from x.x.x.x` changing the x's to your IP.


Image not found or type unknown



You can now type `ip a` on your source system to get its IP you need.




On your source system, now type `python3 -m http.server` to start up a temporary HTTP server on port 8000 for WGet to pull from. If the port is already in use and errors out, type the port after the command, as shown below.




```
root@S01:~# python3 -m http.server 3000
Serving HTTP on 0.0.0.0 port 3000 (http://0.0.0.0:3000/) ..
```

Go to your destination system and go to a folder you want to import into, such as making a directory called `/import` then going into it. Run `wget` `x. x. x. x: PORT/FILENAME. tgz`, replacing the x's with your public or private IP, the PORT

with the port it opened and the FILENAME with the tarball name.



You will see a progress bar with it transferring over. Once it's done, it will go back to the normal terminal line.



You can now go back to the source system and "Ctrl + C" to exit out of the webserver.

```
^C
Keyboard interrupt received, exiting.
root@S01:/var/lib/pterodactyl/volumes#
```

On your destination system, you can now do what you need to do on your tarball, such as move it to another place, extract it with `tar -xvf FILENAME.tgz` and more.

```
031d5d21-2d7a-4381-81ae-3b4602a14985/hub.tgz
root@S02:/home/server# ls
031d5d21-2d7a-4381-81ae-3b4602a14985  hub.tgz
root@S02:/home/server#
```

If the transfer was successful and the extraction also was successful, make sure you go onto the old source system and delete the old tarball, you don't want to fill up on disk space.


```
031d5d21-2d7a-4381-81ae-3b4602a14985
b45  hub.tgz
1d5
umes# rm hub.tgz
umes#
```

Cannot Access UDP Connected Games Due To OVH Firewall


If you have purchased an OVH system from one of their "GAME" series, such as the GAME-1, GAME-2 or the GAME-LE-x servers, you will need to login to their interface to disable the UDP firewall called the "Game Firewall" as this will block connections such as Minecraft Geyser.


This guide presumes that you've checked multiple port open sites and they all state closed, you checked and confirmed the port is open on the firewall internally on UFW or the IPTables rules and if running the Pterodactyl panel, you've confirmed that the port is assigned to the server itself and rebooted the container to confirm working.

Login to your OVH account and click the "Bare Metal Cloud" in the top left, click "Network" along the left hand sidebar then "IP" under the dropdown.




In the list, find your server IP and click the three dots on the right and click "Configure the GAME firewall". If you are unsure which server it is, click the blue text of the IP and it will jump you to that server to confirm it is the one you want to configure.






All you now need to do is click "Deactivating the restriction" along the top bar.



Wait about 2-5 minutes and it should be off. You are now able to connect to that IP with ports broadcasting UDP traffic, like Minecraft Geyser. If it still doesn't work, do it again, as in enabling and disabling it. It's super temperamental and needs it to happen more than 60% of the time!


Accessing A Website Or Control Panel Shows "Too Many Redirects" Error

If you see this error:




Go to CloudFlare and click your domain.

Image not found or type unknown




On the left hand side, expand SSL/TLS then click on "Overview".

Image not found or type unknown



Set the SSL/TLS encryption mode to "Full" instead of "Flexible".


Image not found or type unknown



Wait about 30 seconds for the change to propagate, then you should be able to access it again.

Error 1014 Tebex - CNAME Cross-User Banned

If you are ever presented with the following below when trying to access your Tebex store using your own custom domain, such as <https://store.vanillaplus.uk>, this means you need to **renewal your Plus plan**.



You haven't been banned or suspended by CloudFlare, even though the word "banned" appears. What this means in simple terms is when Tebex tries to redirect from your custom domain to the actual store link, they stop that redirection because you haven't renewed your "Plus" plan, so the CNAME becomes "Banned" to stop the "hop" from your custom domain to *abc123abc123.tebex.io* (The CNAME you made).

Hopefully this provides more context. All you need to do is go and pay. After about 5-10 minutes, the error disappears.

Why Is X Not Installed On The System - How Docker Images Work

When connecting to your systems via SSH, you may notice that certain applications don't exist, are not installed or not working as they should be. This is due to how docker works and it's image system.

If you are a client who has Pterodactyl installed, servers are managed in containers, which are managed by the software called "Docker". Each server is contained in a container, simply referred to as a "house." Those little houses live together to run anything you wish to in their own little "building". They can't talk to other "houses" (*Unless internal, 172.18.0.1 or your own subnet*), they can't interact with them, they can't touch them, they are **separate**. They can run a magnitude of software such as Java, NodeJS, Python, SteamCMD, the list goes on. Now, these are managed with **docker images**. Images are released by many users online to help streamline the process of running a certain application. For example, let's take a java one from Pterodactyl's repository on GHCR:

- ghcr.io/pterodactyl/yolks:java_17

What we can see here is that it's using the website **ghcr.io**, which is where people can share and publish docker images. Next up is **/pterodactyl**, which is the author of the images, then **java_17**, meaning that Java 17 will be installed into that **container** when deployed. Deploying means that you've gone to your panel, created a server, selected Java 17 as the image, and the "house" gets "built." Using this analogy should hopefully now make sense. **Nothing** is installed on the **root** level, **doesn't touch** the system whatsoever and doesn't interfere with **other** houses.

This means that when you run **java --version**, **node -v** or other such command, it won't show anything up unless we installed that application at some point for testing

or development. Any questions, ask the team.

PHP Extensions

Locations For .so Files

PHP stores its .so extension files in /usr/lib/php/ then the API version of PHP. For ease of use, they have been listed here as they are not listed on any site.

- **PHP 8.0:** 20200930
- **PHP 8.1:** 20210902
- **PHP 8.2:** 20220829

If you need it and you can't find it, run **php -i | grep API** and you will see the API number:

```
Server API => Command Line  
PHP API => 20210902  
Zend Extension Build =>
```

Installing Modules On cPanel For WHMCS - IonCube, FileInfo & Soap

IonCube

- Go to [/opt/cpanel/ea-php81/root/usr/lib64/php/modules](#) (**PHP 8.1**) and unzip the tar you download from [here](#) with the .so file. (The loaders from their website)
- Then in [/opt/cpanel/ea-php81/root/etc/php.d](#) (**PHP 8.1**) put **nano 00-ioncube.ini** and put in [zend_extension = /opt/cpanel/ea-php81/root/usr/lib64/php/modules/ioncube_loader_lin_8.1.so](#)
- Once done, you need to **systemctl restart ea-php81-php-fpm**.

Soap

- Copy the **soap.so** file from a normal PHP-8.1 install in the location [/usr/lib/php/20210902](#) (**PHP 8.1**) and put it into [/opt/cpanel/ea-php81/root/usr/lib64/php/modules](#) (**PHP 8.1**).
- Then in [/opt/cpanel/ea-php81/root/etc/php.d](#) (**PHP 8.1**) put **nano 00-soap.ini** and put in [extension=soap.so](#)
- Once done, you need to **systemctl restart ea-php81-php-fpm**.

FileInfo

- Login to WHM (Not cPanel client side, WHM's admin side).
- Search for "Software" along the left side and find "EasyApache4".
- Click the button that says Customize next to Currently Installed Packages.
- Click PHP Extensions along the left hand side in blue (3rd option down).
- Search for FileInfo in the search bar in the top left and click the toggle on the PHP-8.1 version.

- Click Review.
- After looking over the proposed changes, scroll to the bottom of the screen and click provision.
- Once done, you need to **systemctl restart ea-php81-php-fpm** in terminal.

Referenced from: <https://support.cpanel.net/hc/en-us/articles/1500000107181-How-do-I-install-Fileinfo->

If you wish to install extensions for a user of cPanel but don't want to install these extensions globally, follow:

- **Example below is for soap extension**
- Go into that user with su in terminal or use the cPanel's terminal page.
- Run **pear**, this will give you results. Run **pear install Archive_Tar** to install tar.
- Run **pear install PEAR** to install Pear required for SOAP to install. ***Ironic. Lmfao.***
- Run **pear install soap** then install using the channel link.
- **Note:** You might need to use **--force** on some of the commands as the tar version of 1.4.14 breaks soap install.

Note: On all instances above, you can restart Apache's FPM in the WHM interface instead of the command by searching "Services" in the top left then finding "Apache FPM".

LiteSpeed WebServer With Pterodactyl (CyberPanel)

Setup main domain then “panel” as another subdomain. During the subdomain creation, make the location of the files **/panel/public**. Pull files from this directory back one into **/panel** and make sure Ptero's **/public** is the folder that CyberPanel is pointing to for it's VHost - This tricks CyberPanel into loading it.

Remember it's a laravel app, so you still need to install composer, run the composer install, sort out the database connections, fix the initial permissions from the Ptero docs, install redis and sort the pteroq service (including wings if adding it onto this system).

If you get error 500, click fix permissions in file location **/panel** and **/panel/public** in the CyberPanel files interface along the top.

Set **.htaccess** rules in the rewrite section of CyberPanel with:

```
RewriteEngine On
RewriteBase /
RewriteRule ^/index.php$ - [L]
RewriteCond %{REQUEST_FILENAME} !-f
RewriteCond %{REQUEST_FILENAME} !-d
RewriteRule . /index.php [L]
```

Note: If not added, the panel will just 404 as it cannot redirect .php files to friendly URL's.

For wings, it will fail since firewalld is blocking the ports. Go and open the ports in **/etc/firewalld/zones/public.xml** and update the public zone. Then **systemctl**

restart firewall. You can also do this in the Firewall page of CyberPanel along the left hand side.

Searching Through Files

- Finding A Particular String And Updating It

Main command: `grep -rnw "xA" -e "xB" | xargs sed -i 's/xC/xD/'`

- **grep** – Obvious, search to find results.
- **-rnw** – Search all directories after string definition.
- **"xA"** – Location to search.
- **"xB"** – String to search. **Has to be in quotations!**
- **|** – Pipe, put results from grep into sed.
- **xargs** – Call on argument from result.
- **sed** – Replace.
- **'s/xC/xD/'** – Replace string xC with xD.

If you see on the example below, you would use the grep command up to the pipe symbol first without **-l**, then add it in afterwards. This is used **after** searching for strings. The letter **L** is for grep'ing just the directory location where it finds the result and not actually showing file contents.

SED WILL NOT work if you do not add **-l** into the **-rlnw**. It has to require the absolute paths and not anything else it outputs.

```

root@ubuntu:/var/lib/pterodactyl/volumes# grep -rlnw */plugins/*/*.yaml -e '185
73640091-16fa-44b0-a0ec-97c0858bf3d1/plugins/Bounty/config.yaml:2: uri: "mongodb:
min"
73640091-16fa-44b0-a0ec-97c0858bf3d1/plugins/License-System-Tester/config.yaml:1:server: "http
73640091-16fa-44b0-a0ec-97c0858bf3d1/plugins/LifeSteal-Core/config.yaml:3: uri: "mongodb://mong
hSource=admin"
7f19c50c-10de-4801-b797-52a0da1469e2/plugins/LifeSteal-Core/config.yaml:3: uri: "mongodb://mongoe
min"
7f19c50c-10de-4801-b797-52a0da1469e2/plugins/WheelSpin/config.yaml:2: ip: 185.
94f2f2ec-0b6d-4b04-bb72-e99d2e1e8fa3/plugins/Kits/config.yaml:2: ip: 185
94f2f2ec-0b6d-4b04-bb72-e99d2e1e8fa3/plugins/OpenAudioMc/data.yaml:11: server-ip: 185
9c010c0c-3b48-4da3-896b-cf7beb70cb02/plugins/LifeSteal-Core/config.yaml:3: uri: "mongodb://
/?authSource=ad
min"
9c010c0c-3b48-4da3-896b-cf7beb70cb02/plugins/WheelSpin/config.yaml:2: ip: 185.
bd8341c6-0a3e-4ff3-87c9-255130402777/plugins/Bounty/config.yaml:2: uri: "mongodb:
/?authSource=ad
min"
bd8341c6-0a3e-4ff3-87c9-255130402777/plugins/LifeSteal-Core/config.yaml:3: uri: "mongodb:
?aut
hSource=admin"
bd8341c6-0a3e-4ff3-87c9-255130402777/plugins/WheelSpin/config.yaml:2: ip: 185.
dc963007-dddb-4efd-a6ee-b41dbbb7f507/plugins/LifeSteal-Core/config.yaml:3: uri:
'aut
hSource=admin"
ec7551c5-f876-4318-9856-71f183884ab1/plugins/OpenAudioMc/data.yaml:11: server-ip: 185
root@ubuntu:/var/lib/pterodactyl/volumes# grep -rlnw */plugins/*/*.yaml -e '185
73640091-16fa-44b0-a0ec-97c0858bf3d1/plugins/Bounty/config.yaml
73640091-16fa-44b0-a0ec-97c0858bf3d1/plugins/License-System-Tester/config.yaml
73640091-16fa-44b0-a0ec-97c0858bf3d1/plugins/LifeSteal-Core/config.yaml
7f19c50c-10de-4801-b797-52a0da1469e2/plugins/LifeSteal-Core/config.yaml
7f19c50c-10de-4801-b797-52a0da1469e2/plugins/WheelSpin/config.yaml
94f2f2ec-0b6d-4b04-bb72-e99d2e1e8fa3/plugins/Kits/config.yaml
94f2f2ec-0b6d-4b04-bb72-e99d2e1e8fa3/plugins/OpenAudioMc/data.yaml
9c010c0c-3b48-4da3-896b-cf7beb70cb02/plugins/LifeSteal-Core/config.yaml
9c010c0c-3b48-4da3-896b-cf7beb70cb02/plugins/WheelSpin/config.yaml
bd8341c6-0a3e-4ff3-87c9-255130402777/plugins/Bounty/config.yaml
bd8341c6-0a3e-4ff3-87c9-255130402777/plugins/LifeSteal-Core/config.yaml
bd8341c6-0a3e-4ff3-87c9-255130402777/plugins/WheelSpin/config.yaml
dc963007-dddb-4efd-a6ee-b41dbbb7f507/plugins/LifeSteal-Core/config.yaml
ec7551c5-f876-4318-9856-71f183884ab1/plugins/OpenAudioMc/data.yaml
root@ubuntu:/var/lib/pterodactyl/volumes#

```

An Example

Searching through LuckPerms configs in all servers to replace connection timeout string.

```
grep -rlnw "connection-timeout: 5000" */plugins/LuckPerms | xargs sed -i
's/connection-timeout: 5000/connection-timeout: 20000/'
```

Installing SQLite3 To SQL with Python3 Package To Convert SQL-Lite DB to MySQL for import

Start off with **pip3 install sqlite3-to-mysql**, making sure **apt install python3-pip** is installed. Once done, upgrade pip3 with **pip3 install --upgrade pip** then **pip3 install packaging**. You can now run **sqlite3mysql --help**.

For simple syntax, it is:

```
sqlite3mysql -f <file-of-sqlitedb.db> -d <database-to-import-to> -u <sql-  
user-to-import-with> -p
```

When I tried using `-u root`, it didn't work and said access denied even with correct password. Use a SQL user which has access to that particular database or a global admin. *No idea why.*

Works fairly well and have had no reports of issues.

An example below:


```
root@s01 /var/lib/pterodactyl/volumes/24e61f3-e9ea-4fe8-9330-b5b313b1894c/plugins/VotingPlugin # sqlite3mysql -f Users.db -d s1_votingplugin -u blockdb -p MySQL password:  
2023-11-25 01:59:26 INFO Transferring table Users  
100%|██████████████████████████████████████████████████████████████████████████| 3235/3235 [00:00<00:00, 3584.05it/s]  
2023-11-25 01:59:27 INFO Done!
```

ItemsAdder With CloudFlare Proxied Domain For IP Protection

- Create a proxy pass for the domain and add to CloudFlare.
- Cert it with LetsEncrypt and give the port to the server on the panel, like 20000.
- Update ItemsAdder with the details below and make all the other "host" types false.
- This texture pack option should fully deploy automatically, and the game should recognise it.
- If it doesn't, open the MC log and see if it "Times out", if it does, you've done a configuration wrong!


```
self-host:  
  enabled: true  
  server-ip: https://texturepack.domain.com  
  pack-port: 20000  
  append-port: false
```

image not found or type unknown





Get the “/dev/disk/by-id” information for FSTab from the udevadm command

To start with, you will need to use the disk name **only** without the partition name. When you go into */dev/disk/by-id* and investigate the names, it will be named “part-1, part-2” etc at the end of the ID name. Ignore these for now. Do *fdisk -l* and get the disk name.



Run the below command:

```
udevadm info -q symlink --path=/sys/block/<DRIVE-NAME> | awk '{print "/dev/" $1}'
```



Repeat this command until you have “**nvme-eui.ID**” or “**sci-ID**” as the first time you run it, it may list the “**by-path/**” statement which isn't helpful or the actual “**by-id/DISKNAME**” which is what we want.

Once you've found the ID, go to [/dev/disk/by-id](#) and look for the partition you want from [fdisk -l](#). Use this to correlate the lists together.

Device	Start	End	Sectors	Size	Type
/dev/sda1	2048	4095	2048	1M	BIOS boot
/dev/sda2	4096	4198399	4194304	2G	Linux filesystem
/dev/sda3	4198400	41023266815	41019068416	19.1T	Linux filesystem

```
root@node1:/dev/disk/by-id# ls
dm-name-ubuntu--vg-ubuntu--lv
dm-uuid-LVM-0x2FJev1XxGIiK3hSYv9017LUdAdPuC0C7WnLnUXTD1eieg1Rq1WLxNTCcUsThAv
lvm-pv-uuid-cLhLmV-G8rW-mN0c-ly8B-zqL0-AGkF-XYPJLc
nvme-CT4000P3PSSD8_2323E6DFEEFA
nvme-CT4000P3PSSD8_2323E6DFEEFA-part1
nvme-CT4000P3PSSD8_2323E6DFEF05
nvme-CT4000P3PSSD8_2323E6DFEF05-part1
nvme-eui.6479a77bf000006d
nvme-eui.6479a77bf000006d-part1 ←
nvme-eui.6479a77bf000009b
nvme-eui.6479a77bf000009b-part1 ←
scsi-331402ec0154ab8d0
scsi-331402ec0154ab8d0-part1 ←
scsi-331402ec0154ab8d0-part2 ←
scsi-331402ec0154ab8d0-part3 ←
scsi-SATA_CT1000MX500SSD1_2321E6DBA73F
scsi-SATA_CT1000MX500SSD1_2321E6DBA73F-part1
scsi-SATA_CT1000MX500SSD1_2321E6DBA73F-part2
scsi-SATA_CT1000MX500SSD1_2321E6DBA73F-part3
wwn-0x31402ec0154ab8d0
wwn-0x31402ec0154ab8d0-part1
wwn-0x31402ec0154ab8d0-part2
wwn-0x31402ec0154ab8d0-part3
```

For example, [/dev/sda3](#) is going to be **"part-3"** from the ID list.

```
Disk /dev/nvme0n1: 3.64 TiB, 4000787030016 bytes, 7814037168 sectors
Disk model: CT4000P3PSSD8
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: gpt
Disk identifier: E939BE4A-F338-A847-BF08-15941DFD4C6F

Device      Start      End      Sectors  Size Type
/dev/nvme0n1p1 2048 7814037134 7814035087 3.6T Linux filesystem Part 1

Disk /dev/nvme1n1: 3.64 TiB, 4000787030016 bytes, 7814037168 sectors
Disk model: CT4000P3PSSD8
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: gpt
Disk identifier: EFF4E460-A8D8-534D-9231-7677EA87C983

Device      Start      End      Sectors  Size Type
/dev/nvme1n1p1 2048 7814037134 7814035087 3.6T Linux filesystem Part 1

Disk /dev/sda: 931.51 GiB, 1000204886016 bytes, 1953525168 sectors
Disk model: CT1000MX500SSD1
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 4096 bytes
I/O size (minimum/optimal): 4096 bytes / 4096 bytes
Disklabel type: gpt
Disk identifier: 484598C5-0B95-44E5-86D2-305FBA1E4BF8

Device      Start      End      Sectors  Size Type
/dev/sda1    2048      2203647 2201600    1G EFI System Part 1
/dev/sda2    2203648    6397951 4194304    2G Linux filesystem Part 2
/dev/sda3    6397952   1953521663 1947123712 928.5G Linux filesystem Part 3
```

DO NOT IGNORE THIS WARNING. IT IS CRUCIAL TO READ.

Note: Make sure you set the drive to ext4 before rebooting so it actually knows how to handle the partition. If you don't, the partition won't actually mount and when you reboot you will get put into **initramfs**, so have rescue mode on standby with a KVM!

```
root@node1:/dev/disk/by-id# mkfs.ext4 /dev/nvme0n1p1
```

Go into nano [/etc/fstab](#) and add the [/dev/disk/by-id/IDNAME-PARTX](#) into the FSTab. Example below for **nvme0n1p1's part 1, the partition for the datastore, ext4, defaults, 0, 1.**

```
# NVMe Drives
/dev/disk/by-id/nvme-eui.6479a77bf000006d-part1 /srv/datastore1 ext4 defaults 0 1
/dev/disk/by-id/nvme-eui.6479a77bf000009b-part1 /srv/datastore2 ext4 defaults 0 1
```


Resizing single volume group and logical volume after disk expansion

This guide is used after cloning to a increased size disk such as 500GB to 1TB, or by having some space left over from installing the OS, such as the rest of the disk but only 50GB allocated to the / partition.

Firstly run **vgdisplay** to get the volume group information.

```
--- Volume group ---
VG Name                vg0
System ID
Format                 lvm2
Metadata Areas         1
Metadata Sequence No   12
VG Access               read/write
VG Status               resizable
MAX LV                 0
Cur LV                 2
Open LV                 2
Max PV                 0
Cur PV                 1
Act PV                 1
VG Size                 424.18 GiB
PE Size                 4.00 MiB
Total PE                238073
Alloc PE / Size         238073 / 424.18 GiB
Free PE / Size          0 / 0
VG UUID                1rhMP5-SSbG-7vzf-SRwG-6Uzk-8sX9-ufVTr6
```

Run **lsblk** and notice that the size of the partition on the disk **and** the volume group is 424G, not the 931.5G.

Note: After resizing, it will be a bit smaller than the disk as the other partitions take it up, so around 926G at the end

image not found or type unknown



Let's add the rest of the free space to the partition first. Run **cfdisk /dev/nvme#n1**, replacing hashtag with drive number. No need for a partition here as we need to do changes on the disk itself.

I should have an image here of CF Disk at some point when I have a free size again.

Go down to partition 3 (Or where your logical volumes are stored for expansion) and chose "Resize". Just put in the amount you wish to add to that partition (which is presumably max), then write it.

Image not found or type unknown




Image not found or type unknown




Image not found or type unknown



Once you've done this, run **pvresize /dev/nvme#n1p#**, replacing both the hashtags respectively.

image not found or type unknown



Next run **lvextend -l +100%FREE /dev/vg#/LOCATION**, replacing the hashtag with the volume group number, normally vg0, and location with the mountpoint name, normally "root". An example:


This will say resized. I don't have an image currently of this.

```
root@S01:~# lvextend -l +100%FREE /dev/vg0/root
New size (237049 extents) matches existing size (237049 extents).
root@S01:~#
```

Once done, run **resize2fs /dev/vg#/LOCATION**, replacing the values again with the relevant fields.


This will say "online sizing required" and resize.

image not found or type unknown



Run **df -h** and **lsblk** to confirm resizing.

image not found or type unknown




Resizing a logical group then volume within a parented RAID number

Taken from <https://packetpushers.net/blog/ubuntu-extend-your-default-lvm-space/>

Run **cat /proc/mdstat** - Get the RAID partition you want to update after doing all the resize commands to switch from Raid 1 to Raid 0.

Taken from <https://serverfault.com/questions/915284/is-it-possible-to-convert-raid1-to-raid0-without-system-reinstallation>

Image not found or type unknown



```
md127 : active raid0 nvme1n1p3[2] nvme0n1p3[0]  
973281280 blocks super 1.2 64k chunks
```

Next, get the partition info which should look like this with **lsblk**:

```

root@S01:/dev/vg0# lsblk
NAME                                MAJ:MIN RM   SIZE RO TYPE MOUNTPOINTS
loop0                               7:0      0    62M  1 loop /snap/core20/1587
loop2                               7:2      0   79.9M  1 loop /snap/lxd/22923
loop3                               7:3      0   38.8M  1 loop /snap/snapd/21759
loop4                               7:4      0   63.9M  1 loop /snap/core20/2318
loop5                               7:5      0    87M  1 loop /snap/lxd/28373
nvme1n1                             259:0     0  465.8G  0 disk
├─nvme1n1p1                         259:5     0    550M  0 part /boot/efi
├─nvme1n1p2                         259:6     0     1G  0 part
│   └─md126                         9:126     0     2G  0 raid0 /boot
├─nvme1n1p3                         259:7     0  464.2G  0 part
│   └─md127                         9:127     0  928.2G  0 raid0
│       └─vg0-swap                  253:0     0     4G  0 lvm  [SWAP]
│           └─vg0-root              253:1     0  924.2G  0 lvm  /
nvme0n1                             259:1     0  465.8G  0 disk
├─nvme0n1p1                         259:2     0    550M  0 part
├─nvme0n1p2                         259:3     0     1G  0 part
│   └─md126                         9:126     0     2G  0 raid0 /boot
├─nvme0n1p3                         259:4     0  464.2G  0 part
│   └─md127                         9:127     0  928.2G  0 raid0
│       └─vg0-swap                  253:0     0     4G  0 lvm  [SWAP]
│           └─vg0-root              253:1     0  924.2G  0 lvm  /
root@S01:/dev/vg0#

```

Notice how **vg0-root** and **vg0-swap** are sitting on on both disks **partition 3** parented under RAID mdadm **md127**? This is what we will be working on.

Run `vgdisplay` - Get **V**olume **G**roup name and the path, in this case will be `/dev/vg0/xxx`. Sometimes it is `/dev/mapper/ubuntu-vg-ubuntu-lv`, different on how the OS was installed.

```

root@S01:/dev/vg0# vgdisplay
--- Volume group ---
VG Name                vg0
System ID

```

Increasing the size now comes in 3 steps

1. We now need to physically resize the volume of the mdadm RAID by doing **pvresize /dev/RAID#** as shown below:

```

root@S01:/dev/vg0# pvresize /dev/md127
Physical volume "/dev/md127" changed
1 physical volume(s) resized or updated / 0 physical volume(s) not resized
root@S01:/dev/vg0#

```

2. Expand the **Logical Volume** by running **lvextend -l +100%FREE /dev/VGNAME/LVNAME**. We worked out the **VGNAME** was **vg0** and we worked out the **LVNAME** is **root** from **lsblk** - this showed us **vg0-root**. It should be successful.

```
root@S01:/dev/vg0# lvextend -l +100%FREE /dev/vg0/root
Size of logical volume vg0/root changed from 460.09 GiB (117784 extents) to 924.19 GiB (236593 extents).
Logical volume vg0/root successfully resized.
```


3. Your final stage is doing the typical **resize2fs /dev/VGNAME/LVNAME**. You do **NOT** do it on the mdadm **RAID#**, as this will just fail with not finding the superblock:

```
root@S01:/dev/vg0# resize2fs /dev/md127
resize2fs 1.46.5 (30-Dec-2021)
resize2fs: Device or resource busy while trying to open /dev/md127
Couldn't find valid filesystem superblock.
root@S01:/dev/vg0# resize2fs /dev/vg0/root
resize2fs 1.46.5 (30-Dec-2021)
Filesystem at /dev/vg0/root is mounted on /; on-line resizing required
old_desc_blocks = 58, new_desc_blocks = 116
The filesystem on /dev/vg0/root is now 242271232 (4k) blocks long.
```

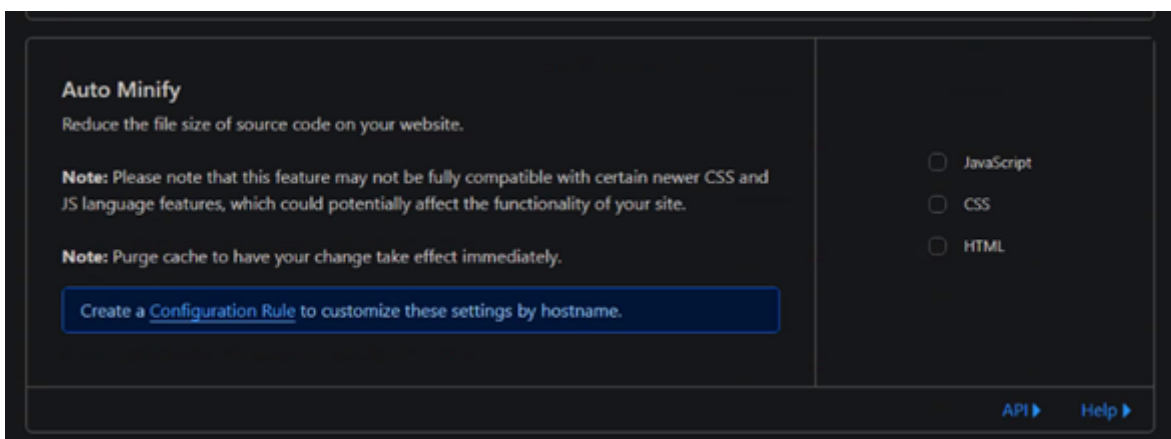
Check with **df -h**, and you will see it has been resized successfully:

```
root@S01:/dev/vg0# df -h
Filesystem      Size  Used Avail Use% Mounted on
tmpfs           6.3G  2.1M  6.3G   1% /run
/dev/mapper/vg0-root 909G  387G  481G  45% /
```



CloudFlare does not minify javascript correctly when working with Grafana



Speed > Optimisation > Content Optimisation.



Follow the below if you want to disable JavaScript Minify'ing for **just that** website, not the whole domain:



Install Custom Version Of MySQL Server

In this guide, replace **.27** with what version you need. At time of writing, Version **.37** is now out, but **on** and **after .29** breaks games such as Unturned with their charsets.

To begin, download the bundle.

```
wget https://dev.mysql.com/get/Downloads/MySQL-8.0/mysql-server_8.0.27-1ubuntu20.04_amd64.deb-bundle.tar
```

Extract it.

```
tar -xf mysql-server_8.0.27-1ubuntu20.04_amd64.deb-bundle.tar
```

Install LibSSL otherwise it will fail.

```
wget http://archive.ubuntu.com/ubuntu/pool/main/o/openssl/libssl1.1_1.1.1f-1ubuntu2_amd64.deb
sudo dpkg -i libssl1.1_1.1.1f-1ubuntu2_amd64.deb
```

Install Libaio1 ([Lo! - AIO](#) [IO](#)) otherwise it will fail.

```
wget http://archive.ubuntu.com/ubuntu/pool/main/liba/libaio/libaio1_0.3.110-5_amd64.deb
sudo dpkg -i libaio1_0.3.110-5_amd64.deb
```

Install MySQL common.

```
dpkg -i mysql-common_8.0.27-1ubuntu20.04_amd64.deb
```

Without the above being done, it errors out stating it needs MySQL common. If you try to include the package (*./mysql-common....*) with *apt install*, this fails due



to it being the "last" package it installs and fails to understand when to install it.

Finally, install MySQL.

```
apt install --yes --no-install-recommends ./mysql-community-client_8.0.27-1ubuntu20.04_amd64.deb ./mysql-community-client-core_8.0.27-1ubuntu20.04_amd64.deb ./mysql-community-client-plugins_8.0.27-1ubuntu20.04_amd64.deb ./mysql-server_8.0.27-1ubuntu20.04_amd64.deb ./mysql-community-server_8.0.27-1ubuntu20.04_amd64.deb ./mysql-client_8.0.27-1ubuntu20.04_amd64.deb ./mysql-community-server-core_8.0.27-1ubuntu20.04_amd64.deb
```

Enter root password for password authentication or leave blank to just run **mysql** and

login with terminal.



Choose "Use Strong Password...", as you won't be needing 5.x compatibility.

Complete. Run **mysql**.

```
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 8
Server version: 8.0.27 MySQL Community Server - GPL

Copyright (c) 2000, 2021, Oracle and/or its affiliates.

Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.

mysql> █
```

Networking With Hetzner

- With and Without additional MAC's configured

IP's have been redacted in this configuration but you can evaluate based on your operations.

With doing this in debian, reverting over to netplan and disabling networking has been so much easier. ChatGPT and other help articles from stack overflow has been a major help in understanding networking as a whole. **systemctl stop networking && systemctl disable networking**. Install netplan.io with **apt install netplan.io**.

Configure it then **reboot** the system. **Do not** presume it will work unless you restart.

If you just need the netplan configuration:

```
network:
  version: 2
  renderer: networkd
  ethernets:
    enp41s0: # Network interface name from "ip a".
      dhcp4: no # Turn off DHCP.
      addresses:
        - 148. x. x. 242/29 # IP Address 1 with subnet.
        - 144. x. x. 121/27 # IP Address 2 with subnet.
```

```

gateway4: 148. x. x. 241 # The gateway of the first subnet.
nameservers:
  addresses:
    - 8.8.4.4 # Obvious enough.
    - 8.8.8.8 # Obvious enough.
routes:
  - to: 144. x. x. 96/27 # The "Identifier IP" for the subnet. The one before the gateway
and two before the first usable.
    via: 144. x. x. 97 # The gateway of this subnet.
    metric: 100 # The higher the metric, the lower it is in priority (Yes, stupid) - It
means if 240 fails, it will make 144 prioritized.

```

If you need to specify MAC addresses:

```

network:
  version: 2
  renderer: networkd
  ethernets:
    enp41s0: # Network interface name from "ip a".
      dhcp4: no # Turn off DHCP.
      addresses:
        - 148. x. x. 242/29 # IP Address 1 with subnet.
        - 144. x. x. 121/27 # IP Address 2 with subnet.
      macaddress: xx:xx:xx:xx:xx:xx # Specify the MAC address for 148. x. x. 242 ONLY (It has to
route somewhere for remaining traffic).
      routes:
        - to: 144. x. x. 96/27 # The "Identifier IP" for the subnet. The one before the gateway
and two before the first usable.
            via: 144. x. x. 97 # The gateway of this subnet.
            metric: 100 # The higher the metric, the lower it is in priority (Yes, stupid) - It
means if 240 fails, it will make 144 prioritized.
        - to: 148. x. x. 240/29 # The "Identifier IP" for the subnet. The one before the gateway
and two before the first usable.
            via: 148. x. x. 241 # The gateway of this subnet.
            metric: 0 # The higher the metric, the lower it is in priority (Yes, stupid) - Not
applicable here - 0 meaning normal priority.

-----
-

# If using multiple interfaces, add them here with their respective MAC addresses

```

enp41s0: 1:

dhcp4: no

addresses:

- 144. x. x. 121/27

macaddress: yy: yy: yy: yy: yy: yy # MAC address for 144. x. x. 121

Running Lancache & Another Webserver On Top

If you are running the software Lancache and need to run another webserver on top such as nginx with ports 80 and 443, you need to do some configuration changes by serving your virtual host files on the public IP address while your Lancache speaks on your private internal IP's to your dedicated systems/LAN network.

In this example, I will use Lancache on 10.10.10.50 and the public IP as 123.123.123.123. What we need to do first is configure Lancache with the following inside the .env file:

```
## IP addresses that the lancache
## Specify one or more IPs, space
## Note: This setting only affects
LANCACHE_IP=10.10.10.50

## IP address on the host that the
DNS_BIND_IP=10.10.10.50
```

Once this has been configured, head into the docker-compose.yml and update the **80** and **443** ports to be bound onto that particular bind of private IP.

```
monolithic:
  image: lancachenet/monolithic:latest
  env_file: .env
  restart: *restart-policy
  ports:
    - "10.10.10.50:80:80"
    - "10.10.10.50:443:443"
```


Once completed, start up Lancache **first** as it needs port 80 and 443 in use for the Monolithic container using **docker-compose up**. Check that it was successful, then **Ctrl+C**, then **docker-compose up -d**.

We've now configured Lancache to work on that private LAN IP address. For all your virtual host configuration files now for nginx, you just need to specify another private LAN IP, or the public IP if serving accessible websites (if you aren't using HAProxy).

Image not found or type unknown





Image not found or type unknown



You can now restart nginx and it will not complain that the port is already in use as you've done this neat hack.

Keep in mind that if you have software that updates the nginx vhost files automatically such as TenantOS or VirtFusion, please keep in mind that you will probably need to repeat this again on their updates.

Installing Debian/Ubuntu & Then Disk Unmounts - Debootstrap Error


You're here because you've landed yourself at the error below, stating that it could not find the codename for the release during an installation of Ubuntu or Debian. This is because the CD unmounted during your ISO install, so after provisioning the disks, it's no longer there and cannot continue. Follow below to fix it

Image not found or type unknown




Press enter/escape and head back onto the installer page, then going to "Execute a shell".

Image not found or type unknown



Press continue.

Image not found or type unknown



You're now in the BusyBox shell. Type "**blkid**" and find the CD drive. It's got the name of the ISO.

Image not found or type unknown



In this example, we can see it has the type **iso9660**, so we know that's the ISO, and the mount is **/dev/sr2**. Now simply mount it! **mount /dev/sr2 /cdrom**.

Make sure it's **/cdrom** as this is what the installer is looking for.

Type exit then press back on install.

Image not found or type unknown



Job done!

