

System Admin

- Audio Devices
- Configure FTP
- Configure Postfix
- Configuring Multi-boot Filesystems
- Mounting Default Filesystems
- Server Hostname
- Swap Allocation
- Synchronizing Time Using NTP
- System Sensors
- Unattended Upgrades
- Crontab
- Systemd Services

Audio Devices

When using various Linux Distributions, you may (or may not) run into some issues with audio devices. See some of the configs, logs, and commands below for helpful output in troubleshooting these issues.

GUI Tools / Applications

If looking for a GUI Tool to select or view output / input audio devices, check out `pavucontrol` -

```
sudo pacman -Syu pavucontrol
pavucontrol
```

will install and open the application, which provides a simple interface for selecting audio devices, and even provides application-level audio control, which enables you to easily specify the devices for individual applications instead of forcing a system-wide audio setting for all running apps.

Commands

:)

Sound Card / Devices

Search for all connected audio cards, and output the result.

```
aplay -L | grep :CARD
```

List all connected PCI devices (Sound cards are a PCI device)

```
lspci
```

Audible Sound Test

The command below will send static to each speaker connected to the device, sequentially, one at a time. Running this will continually test all speakers on a loop, until the user exits with `CTRL+C`.

```
speaker-test -D default:PCH -c 8
```

The output from the above test will look similar to the below, depending on your system and devices.

The `-D` argument specifies the audio device you want to test. This is useful when not entirely sure which device is valid, you can test quickly with this cmd and make changes later in `alsamixer` or another config tool with the results of your findings.

The `-c` argument specifies the number of audio channels you want to test, for my setup I only have a front left and right speaker, so 2 will suffice. If I had a surround sound with Left / Right speakers in the back and an additional center speaker, we would test over 5 channels.

```
[kapper@kapper-pc ~]$ speaker-test -D default:PCH -c 2
```

```
speaker-test 1.1.9
```

```
Playback device is default:PCH
```

```
Stream parameters are 48000Hz, S16_LE, 2 channels
```

```
Using 16 octaves of pink noise
```

```
Rate set to 48000Hz (requested 48000Hz)
```

```
Buffer size range from 2048 to 16384
```

```
Period size range from 1024 to 1024
```

```
Using max buffer size 16384
```

```
Periods = 4
```

```
was set period_size = 1024
```

```
was set buffer_size = 16384
```

```
0 - Front Left
```

```
1 - Front Right
```

```
Time per period = 5.648263
```

```
0 - Front Left
```

```
1 - Front Right
```

```
Time per period = 5.973649
```

```
0 - Front Left
```

```
^CWrite error: -4, Interrupted system call
```

```
xrun_recovery failed: -4, Interrupted system call
```

```
Transfer failed: Interrupted system call
```

Sound Mixer / Settings

To open alsamixer, run the below and use the `F6` key to ensure the proper device is selected. This tool can also be used to change volume levels, be careful messing with settings you are unfamiliar with, you could easily blow a speaker. At the least, connect a cheaper pair.

```
alsamixer
```

To check device audio settings / levels via CMD -

```
amixer
```

 to list devices and settings

```
amixer sset Master unmute
```

 to mute the Master device. Master can be changed to any valid device name given from the output of `amixer`

Also, see [Advanced Linux Sound Architecture](#) for more information on various documented issues encountered.

Configure FTP

You can use `scp` instead, taking advantage of the added security and already-configured users on your system. It works a lot like `ssh`

Copy a file from a remote host to your local host -

```
scp -i ~/.ssh/some_key -P 22 username@123.123.123.12: /home/username/test .
```

If you still need or want FTP, you can follow the steps below to configure the FTP server and then connect with Filezilla.

Installing Very Secure FTP Daemon

I am using an Ubuntu 19.04 server in this guide, depending on your system your steps may vary slightly.

Assuming you have nothing installed, run `sudo apt-get update && sudo apt install vsftpd` to install vsftpd (Very Secure FTP Daemon). Navigate to the home directory of the user you wish to enable FTP access, and run the following.

```
# Login as your sudo user
sudo su USER

# Create FTP Directory
mkdir /home/USER/ftp
sudo chown nobody:nogroup /home/USER/ftp
sudo chmod a-w /home/USER/ftp
```

Create User FTP Directories

Create a directory where files can be uploaded, you can name this directory whatever you want. Give this directory permissions so you can upload files to it via FTP clients like FileZilla.

```
mkdir /home/USER/ftp/files
```

```
sudo chown USER: USER /home/USER/ftp/files
sudo chmod 777 /home/USER/ftp/files
```

Configure vsftpd Settings

If you have a firewall enabled, be sure you open the TCP ports 20, 21, 990, and 40000-50000 before you continue.

Add the following to `/etc/vsftpd.conf`

```
# FTP Initial Configuration Options
pasv_min_port=40000
pasv_max_port=50000
user_sub_token=$USER
local_root=/home/$USER/ftp
pasv_min_port=40000
pasv_max_port=40000
userlist_enable=YES
userlist_file=/etc/vsftpd.userlist
userlist_deny=NO
pasv_promiscuous=YES
allow_anon_ssl=NO
force_local_data_ssl=YES
force_local_logins_ssl=YES
ssl_tlsv1=YES
ssl_sslv2=NO
ssl_sslv3=NO
require_ssl_reuse=NO
ssl_ciphers=HIGH
```

Run `echo "USER" | sudo tee -a /etc/vsftpd.userlist && sudo systemctl restart vsftpd` to add your user to the userlist file we configured above and restart the service

`tail -f /var/log/syslog` in another console to see a live feed of service logs when restarting instead of checking the status with `sudo systemctl status`

Change or modify the following values, when editing these files I like to comment out the default value, and create a separate value in an organized list with my custom settings. This is useful later should I want to refer back to the default value I can just search it up in my file, and keeps things

organized so when I can pick things back up quickly. You can do this however you see fit. The result of my modified values within `vsftpd.conf` is below.

```
# Values Modified During FTP Setup
chroot_local_user=YES
write_enable=YES
ssl_enable=YES
```

Run `sudo systemctl restart vsftpd` to restart the service and test your connection using [Filezilla](#).

Debugging FTP Connections

If you're having issues with your FTP connection, check on the service with the following commands `sudo systemctl -l status vsftpd` `sudo tail -f /var/log/vsftpd.log`

To test FTP connections via commandline, run the following

```
ftp -p IPADDRESS
```

You cannot connect to FTP via commandline using this method if you have enabled SSL/TLS because your connection will *not* be encrypted under TLS. Use Filezilla or another encrypted connection method instead.

Notes

Here's a working config file, with some comments on some extra settings I found on the manpages for `vsftpd.conf`

```
# *Example config file /etc/vsftpd.conf.bak
# *Don't forget to backup your default /etc/vsftpd.conf.bak
#
# Custom FTP configuration for basic server configuration
#
# These settings should be refined for security
# Firewall should be used and reflect the settings in this file
# For more security, use keys and disable password authentication
# +Restrict FTP access to a list of approved IP's with distributed keys

# FTP Custom Configuration Options
```

```
# Set chroot user options
chroot_local_user=YES
user_sub_token=$USER

# Set Directory FTP Will Default Into
local_root=/home/$USER/ftp
write_enable=YES
# If you can't write with Write_enable=YES, check directory permissions
# Create .../ftp/files and chmod 777 .../ftp/files

# Passive FTP Connection Settings
pasv_promiscuous=YES
pasv_min_port=40000
pasv_max_port=50000

# userlist_enable=YES tells vsftpd to read /etc/vsftpd.userlist
# /etc/vsftpd.userlist should contain one user per line
userlist_enable=YES
userlist_file=/etc/vsftpd.userlist
# Sets the userlist to be a whitelist or a blacklist
# userlist_deny=YES will deny FTP for any user on the list
userlist_deny=NO
# Enable logs for failed FTP connections due to userlist errors
# userlist_log=YES

# Enable dual logs for vsftpd in /var/log/
# log/xferlog - standard parsable log
# log/vsftpd.log - vsftpd formatted logs
dual_log_enable=YES

# This option specifies the location of the RSA certificate to use for SSL
# encrypted connections.
rsa_cert_file=/etc/ssl/certs/ssl-cert-snakeoil.pem
rsa_private_key_file=/etc/ssl/private/ssl-cert-snakeoil.key
# Other Settings For SSL
ssl_enable=YES
```

```
allow_anon_ssl=NO
force_local_data_ssl=YES
force_local_logins_ssl=YES
ssl_tlsv1=YES
ssl_sslv2=NO
ssl_sslv3=NO
require_ssl_reuse=NO
ssl_ciphers=HIGH

# Default vsftpd.conf Values

# READ THIS: This example file is NOT an exhaustive list of vsftpd options.
# Please read the vsftpd.conf.5 manual page to get a full idea of vsftpd's
# capabilities.

# The default compiled in settings are fairly paranoid. This sample file
# loosens things up a bit, to make the ftp daemon more usable.
# Please see vsftpd.conf.5 for all compiled in defaults.

secure_chroot_dir=/var/run/vsftpd/empty
pam_service_name=vsftpd
connect_from_port_20=YES
use_localtime=YES
dirmessage_enable=YES
local_enable=YES
anonymous_enable=NO
listen_ipv6=YES
listen=NO
```

Modifying the values below during setup of TLS encryption caused vsftpd to crash on startup.. These values were obtained following [this tutorial](#). Just noting this in case I missed something here, so I can revisit it later.

```
# Working values, establishes TLS connection via Filezilla FTP
rsa_private_key_file=/etc/ssl/private/ssl-cert-snakeoil.key
rsa_cert_file=/etc/ssl/certs/ssl-cert-snakeoil.pem

# Modified values from generating ssl cert that are crashing vsftpd
```

```
# rsa_cert_file=/etc/ssl/private/vsftpd.pem  
# rsa_private_key_file=/etc/ssl/private/vsftpd.pem
```

Configure Postfix

“ Postfix is a Mail Transfer Agent (MTA) that can act as an SMTP server or client to send or receive email. There are many reasons why you would want to configure Postfix to send email using Google Apps and Gmail. One reason is to avoid getting your mail flagged as spam if your current server's IP has been added to a blacklist.

[Linode Postfix Tutorial](#)

Install postfix and mailutils -

```
sudo apt install postfix mailutils
```

Create Google App Token

When attempting to send mail from a new host, you may encounter errors with Google blocking or filtering your mail as spam. To prevent this, simply create a Gmail account you wish to send the mail under, [Activate 2FA](#) on the new account, then [Generate App Tokens](#) to distribute to the hosts / apps you wish to send mail on your behalf. See below for further instructions once you have a Gmail account created, and have generated an app password / token.

Postfix App Token Authentication

Once you have the app token, we'll need to add it to `/etc/postfix/sasl/sasl_passwd` - If this file doesn't already exist, create it and include the following lines, modified with your information

```
sudo echo "[smtp.gmail.com]:587 username@gmail.com: password" > /etc/postfix/sasl/sasl_passwd;
```

Instead of using the password you usually input when logging into the GMail account, add the app token generated after enabling 2FA following the links in the first step above. Below, we notify postfix that we've made these changes by running `sudo postmap /etc/postfix/sasl/sasl_passwd`. This will create a `sasl_passwd.db` file in the `/etc/postfix/sasl` directory.

Run `postmap`, and restrict access to our new file containing this password

```
sudo postmap /etc/postfix/sasl/sasl_passwd;
sudo chown root:root /etc/postfix/sasl/sasl_passwd /etc/postfix/sasl/sasl_passwd.db;
sudo chmod 600 /etc/postfix/sasl/sasl_passwd /etc/postfix/sasl/sasl_passwd.db;
```

Configure Relay Server

Configure postfix to relay mail through GMail's server by making the below changes to

`/etc/postfix/main.cf` -

```
# Change / modify this line.
relayhost = [smtp.gmail.com]:587

# Add these lines...
# Enable SASL authentication
smtp_sasl_auth_enable = yes
# Disallow methods that allow anonymous authentication
smtp_sasl_security_options = noanonymous
# Location of sasl_passwd
smtp_sasl_password_maps = hash:/etc/postfix/sasl/sasl_passwd
# Enable STARTTLS encryption
smtp_tls_security_level = encrypt
# Location of CA certificates
smtp_tls_CAfile = /etc/ssl/certs/ca-certificates.crt
```

Send Mail

That's it! Now restart postfix with `sudo systemctl restart postfix` and test sending mail using any of the commands below -

```
echo "This email confirms that Postfix is working" | mail -s "Testing Postfix"
emailuser@example.com
```

or..

```
sendmail emailaddress@gmail.com
FROM: admin@sub.domain.com
SUBJECT: Hi
Body test text
.
```

Mailer Daemon

To change the email the system sends security alerts to, modify the `/etc/aliases` file to use your email address for the `root` field below. If this isn't already in the file, add it, and run `sudo newaliases` to update the system with the new information.

```
# See man 5 aliases for format
postmaster:    root
root: someone@somedomain.com
```

Now to test that this works correctly, attempt to `sudo` somewhere on the system where you'll be required to enter a password, and botch it - all three times. You'll get an email from your server warning you of the security event! Missing a password on an attempt to `sudo` is a security event

Configuring Multi-boot Filesystems

When installing a fresh Linux Distribution, you might want to dual-boot, or even multi-boot, into different desktop environments. There are some pretty specific requirements we'll need to setup manually for our new partitions though, see below for details on the different partitions needed to setup an open ended multi-boot system alongside windows. This configuration will prompt for selection of OS on boot, and will allow for nearly any number of distributions to be tested alongside each other. These instructions vary slightly based on your specific scenario, so be sure to read and understand the need for each setting below.

Installation Media

When installing *any* operating system, we first need to create our installation media. Sometimes these are distributed as Installation CDs, but a simple USB can be turned into the same thing very easily.

Depending on your system, see the sections below.

Choosing a Distribution

Not sure what distribution to use, or searching for a legit ISO?

[Distrowatch](#) is your friend. They provide rankings, comment boards, forums, and (usually) working links to ISO downloads.

Linux ISO Tools

To burn an image using `umount` and `dd` on Linux, run the commands below.

```
lsblk
```

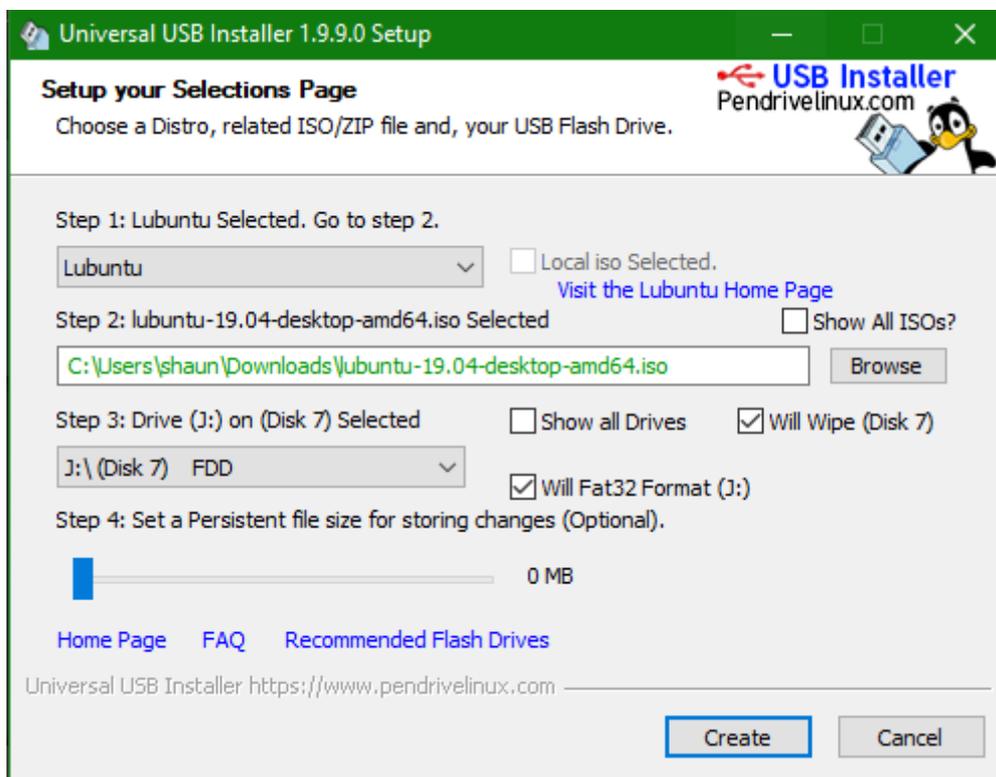
the command above will list all block storage devices connected to your system, find your device in the list, and take note of the name assigned to it within the `/dev/` directory. Usually, this name is `sd<?><?>` or `sd<?>` for the primary partition. Usually, we would want to write to the primary on the USB to ensure that we boot from our ISO.

```
sudo dd if=/home/user/Downloads/inputfile.iso of=/dev/sd<?><?> && sync
```

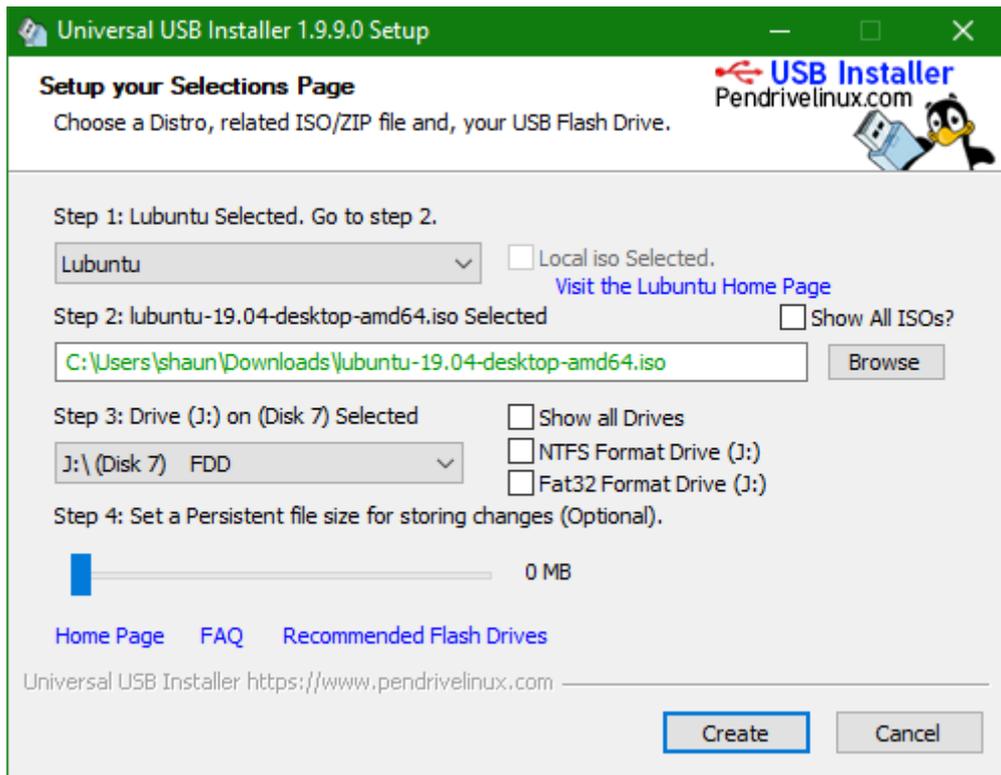
Windows ISO Image Writer

Universal USB Installer handles almost every scenario for most if not all distros. Head over to [UI's Download Page](#), grab the tool and see that your settings are adjusted for your needs.

Your Drive letters, distro release, and persistent file size may be different depending on the requirements you have for your media. To format a drive and prepare for writing -



To create USB installation media -



To create a persistent USB to boot wherever you like, adjust the slider on the bottom of the window labeled **Step 4:** to your desired size.

When using this tool to create a persistent USB device, there is a temporary directory created in your C:\Users\shaun\AppData\Local\Temp\ directory! If there is not enough space on your drive, the process will not terminate itself, but it will not be able to complete. The temporary files tend to be named nsf9D50.tmp or similar, and will take up equal or slightly more space than is being written to your USB. So, if you create a persistent USB with 10GBs of storage on top of writing a 2GB .iso, you can expect to need ~12GB on your C:\ drive in order for the process to complete successfully. Once completed, the temp files should be automatically deleted. If the process gets hung due to insufficient space on your drive, this may not happen and you may need to check your Temp directory to manually delete the files yourself.

Alternatively, you can check out the tools below if you have issues using Universal USB Installer.

[win32diskimager](#)

If you've written to a USB and need to recreate the media for any reason, you'll need to clear the

contents on the USB before you can attempt to reformat and reimage the storage device. To do this, ROSA can be used to clear the USB by selecting the device and clicking clear. After doing this you'll just need to format the drive using windows quick format tool via right click->format->FAT. Then, use win32diskimager to copy the new image to your device.

BIOS Configuration

When configuring a multi-boot system with specific partitions for different distributions, you'll need to enable the following settings within your BIOS -

- Disable Secure Boot
- Disable fast boot
- Set to UEFI mode only within boot options.

If you're unsure how to modify these settings, try running the setting in question through google along with the model of your motherboard. This will hopefully provide some more specific instructions on using the BIOS of your system.

Modifying these settings will allow us to create EFI files within a given EFI partition, created below, where the system defines the boot sequence for multiple operating systems. This allows us to leave our boot sequence open-ended, and easily append EFI system files to our partition / boot options during the installation of a new system. There are, unfortunately, a few discrepancies to how these steps will be performed - unless the system is to be configured *exactly* the same.

USB Boot / Install

Insert your USB installer you created above using [ROSA Image Writer](#) or `dd` command above, and reboot the system. Be sure to pay attention and press the required key to enter the BIOS during boot. For me, the key was delete or F2. Once in the BIOS, navigate to your boot sequence / options and there should be a list of connected storage devices, including all HDD, SSD, USBs, etc. Find your USB installer in the list and select it, this will boot into the installer for your distribution. The installer is usually found on the desktop as an executable application. These installers are usually usable systems, but be aware that there will be no persistent data between reboots until the

installation is completed.

When selecting your installation media to boot from within BIOS, be sure to select the media that corresponds with how your system is configured to boot. In this example, the media should start similar to `UEFI: USB` If you were not using a UEFI configuration, simply select the same media without the `UEFI:` prefix.

Partitioning

Once booted into the USB created above, you will likely see an installation prompt. When given the choice, select 'custom installation or 'custom partition configuration option, and continue with the guide below.

Bootloader Partition

If you are already using Grub on an existing EFI partition, you won't need to create a new one. Skip this step, but make note of where this partition is, we will need it during installation.

This is the partition where we will create and store new bootloaders during installation of different distributions. You will not directly edit or view this partition's contents, but it is the backbone of the system-selection prompt (grub) that you will receive when booting after completing this configuration. There may be a need to step into this partition if you decide to customize your grub configuration, but we won't get into that here.

Size: 1GB (this is generous)

Type: FAT32 **Location:** Beginning of Space (Volume we are partitioning)

Mount: (Leave empty / blank) **Flags:** boot, efi (also called ESP or EFISYS)

You should always choose to install the bootloader on the same disk the EFI filesystem exists, whether your case required the creation of a new EFI volume or if you are installing alongside a previous one. Failing to do so can cause issues during installation.

The only exception to this is when initially installing a Linux / Grub Bootloader - you will have to create a new EFI partition for the Grub Bootloader. Grub will pick up the windows partition

automatically, but if it doesn't, you can always run `sudo grub-update` to search for new EFI partitions or configurations and update your Grub Bootloader appropriately.

Root Partition

This partition will store the Linux system files for your distribution, and unless otherwise partitioned separately, your user's home directory and all of its content. This should be set according to both your distributions total installation size, and if you are not partitioning dedicated space - you should figure in any extra space your user(s) might require for new packages, updates, and applications. Running out of space is a lot worse than having too much, so try to be a little generous here.

Size: Adjust according to installed size of distribution we are using.

Type: ext4 (Logical)

Location: Beginning of Space (Volume we are partitioning)

Mount: / **Flags:** root

Swap Partition (Optional)

This is the space your system will use if you run out of memory. If you max out your RAM, this will prevent your system from freezing up. Be cautious of low RAM systems with little or no swap, the downfall to swap space is that once it is used it cannot be reallocated until the system reboots.

Size: 2GB-Preference (Ideally 50-100% of system RAM)

Type: linuxswap (Logical) **Location:** Beginning of Space (Volume we are partitioning)

Mount: (Leave empty / blank)

Home Partition (Optional)

This is optional. I would recommend having a separate storage device (Massive HDDs are getting cheaper..) to mount your home directories in, so if you ever need to reinstall the root directory of your distributions you'll be able to do so without having to worry about backing up or losing data.

I would not advise taking the gamble, you will probably need to reinstall at some point - and it's good insurance to have.

Size: Preference

Type: ext4 (Primary) **Location:** Beginning of Space (Volume we are partitioning)

Mount: `/home`

Installing

Now all we need to do is specify where to install the bootloader. This is easy since we just created that partition above, the EFI Partition. Select the partition from the dropdown and click 'continue' or 'install' at the bottom corner. After this is complete, you'll just need to reboot and witness the grub! From now on, you'll have an option of which system to boot into when starting your PC.

Adding New Systems

During installation of additional systems, we have two requirements, selecting a location for booting the system, and selecting a location for configuring the root filesystem.

Create a new `/root` (and `/home`, if you choose) partition(s), then select the EFI partition we created above for the bootloader install location (For me, this was sdb1 - the first partition of `/dev/sdb`) The basic requirements of both can be seen below

Bootloader

Size: 1GB (this is generous)

Type: FAT32 **Location:** Beginning of Space (Volume we are partitioning)

Mount: (Leave empty / blank) **Flags:** boot, efi

Root

Size: Adjust according to installed size of distribution we are using.

Type: exfat (Logical) **Location:** Beginning of Space (Volume we are partitioning)

Mount: / **Flags:** root

Grub Issues

If you're having issues with **system options not appearing in grub**, be sure to load into a previous system and run `sudo update-grub` - this command will search for new entries in the EFI partition and automatically add them to your grub configuration / system prompt. You can manually step through the EFI partition using the grub command-line to bail yourself out, but this shouldn't be needed as returning to an already configured system and running this command will pick up all new systems for next reboot.

```
grub rescue> set prefix=(hd0,1)/boot/grub
grub rescue> set root=(hd0,1)
grub rescue> insmod normal
grub rescue> normal
grub rescue> insmod linux
grub rescue> linux /boot/vmlinuz-3.13.0-29-generic root=/dev/sda1
grub rescue> initrd /boot/initrd.img-3.13.0-29-generic
grub rescue> boot
```

```
sudo apt install efibootmgr

sudo modprobe efivars

sudo efibootmgr

sudo efibootmgr -b 4 -B
```

```
test -d /sys/firmware/efi && echo UEFI || echo BIOS
```

```
sudo blkid  
sudo parted -l  
cat /proc/cmdline
```

[Rescue Grub](#)

[Manjaro Install Forum Guide](#)

Mounting Default Filesystems

Surely at some point you will come across the need to mount an additional partition, either temporarily or with persistence across reboots. To do this, you will need to modify the settings within the `/etc/fstab` configuration file. See the details below for more information.

[ArchWiki - Fstab](#)

Fstab Parameters

At first glance, knowing misconfiguration can cause system failure, editing this file can be intimidating. To ensure no mistakes are made, understand what you are specifying when creating a new Fstab entry. Its important to always backup your settings before modifying this file, as in the event of system failure you can always simply restore your misconfiguration with this last working Fstab entry.

From left to right, fstab requires six total parameters, including the UUID. Below, we will go over some settings that might be useful for common configurations.

Device

Describes the block special device or remote filesystem to be mounted

Dir

Describes the mount directory

Type

Sets the file system type

Options

Sets associated mount options

Dump

`dump(8)`

Option to check the filesystem with the `dump(8)` utility. This field is usually set to 0, which disables the check.

“ The fifth field, (`fs_freq`), is used for these filesystems by the `dump(8)` command to determine which filesystems need to be dumped. If the fifth field is not present, a value of zero is returned and `dump` will assume that the filesystem does not need to be dumped.

[man 5 fstab](#)

Fsck

`fsck(8)` Sets the order for filesystem checks at boot time

For the root device it should be 1. For other partitions it should be 2, or 0 to disable checking.

“ The sixth field, (`fs_passno`), is used by the `fsck(8)` program to determine the order in which filesystem checks are done at reboot time. The root filesystem should be specified with a `fs_passno` of 1, and other filesystems should have a `fs_passno` of 2. Filesystems within a drive will be checked sequentially, but filesystems on different drives will be checked at the same time to utilize parallelism available in the hardware. If the sixth field is not present or zero, a value of zero is returned and `fsck` will assume that the filesystem does not need to be checked.

Fstab Entry Examples

The sections below contain example configurations for various fstab scenarios

Mounting Swap Space

For example, to add a swap partition entry created by following the [Swap Allocation](#) instructions, an entry similar to the below would be added to our `/etc/fstab`

```
# <device> <dir> <type> <options> <dump> <fsck>
LABEL=SYSTEM / ext4 defaults 0 0
LABEL=UEFI /boot/efi vfat defaults 0 0
/swapfile none swap defaults 0 0
```

Note that the first line is a comment I tossed into my `/etc/fstab` and is not there by default. The `SYSTEM` and `UEFI` lines are default mount configurations for our root filesystem and boot partitions, respectively. If you have defined a dedicated partition for your `/home` directory, it would also be listed here, sometimes labeled as `DATA`.

Our swap partition is defined by the final line -

```
/swapfile none swap defaults 0 0
```

`/swapfile` points to the swapfile we created following the [Swap Allocation](#) instructions.

`none` tells our system not to mount this device anywhere within our filesystem for access / reading by the user.

`swap` defines the partition as the swap filesystem type

`defaults` passes the default set of options for this fstab entry

`dump` is set to 0, which disables `dump(8)` from running its filesystem check

`fsck` is set to 0, which disables the `fsck(8)` integrity check, since this is swap space and not an

actually filesystem.

Optional Mounts

Sometimes, a filesystem may be removed from a system temporarily. We still use `fstab` to mount this filesystem, but we need to prepare the system for the case that it is not present for mounting at system boot.

“ External devices that are to be mounted when present but ignored if absent may require the `nofail` option. This prevents errors being reported at boot. For example:

```
/dev/sdg1      /media/backup  jfs    nofail, x-systemd.device-  
timeout=1ms   0 2
```

The `nofail` option is best combined with the `x-systemd.device-timeout` option. This is because the default device timeout is 90 seconds, so a disconnected external device with only `nofail` will make your boot take 90 seconds longer, unless you reconfigure the timeout as shown. Make sure not to set the timeout to 0, as this translates to infinite timeout.

[ArchWiki - External Devices](#)

Network Shares

“ If your external device requires another systemd unit to be loaded (for example the network for a network share) you can use `x-systemd.requires=x` combined with `x-systemd.automount` to postpone automounting until after the unit is available. For example:

```
//host/share /net/share cifs noauto,nofail,x-
systemd. automount, x-systemd.requires=network-online.target, x-
systemd.device-
timeout=10,workgroup=workgroup,credentials=/foo/credentials 0 0
```

[ArchWiki - External Devices](#)

Block Device Commands

When creating a new entry within fstab, sometimes you may need to find more information on the devices connected to your system. See the commands below for useful output. None of these commands, when ran verbatim, will make any changes to your system. They are only for gathering needed information.

Note that the output of the commands is present, but has been reduced (and obfuscated, where needed) to keep the length of this page down.

Print information on all connected block devices -

```
lsblk
sda      8:0    0 931.5G  0 disk
├─sda1   8:1    0  128M  0 part
├─sda2   8:2    0 925.5G  0 part
└─sda3   8:3    0   5.9G  0 part
```

Print the UUIDs of all connected block devices, along with some other hardware information -

```
sudo blkid
/dev/sdb2: UUID="436b3ae3-4301-4b8a-80d3-fdf52c7d7059" TYPE="swap" PARTUUID="590670f6-3b89-41b5-b474-fcd6c048628d"
```

Print information on partitions on all connected block devices -

```
sudo parted -l
Model: HDD A12345678-B3210 (scsi)
Disk /dev/sdb: 2000GB
```

Sector size (logical/physical): 512B/4096B

Partition Table: gpt

Disk Flags:

Number	Start	End	Size	File system	Name	Flags
1	1049kB	1000MB	999MB	fat32		boot, esp
2	1000MB	17.0GB	16.0GB	linux-swap(v1)		
3	17.0GB	117GB	100GB	ext4		
4	117GB	217GB	100GB	ext4		
5	217GB	427GB	210GB	ext4		

Print information given a specific block device (partitions) -

```
sudo tune2fs -l /dev/sdb3
tune2fs 1.45.4 (23-Sep-2019)
Filesystem volume name: <none>
Last mounted on: /
Filesystem UUID: fagbraetd325t9-6gafdee7-4d2344agdd-93d2-6f4safsafsa5d6
Filesystem magic number: 0xEF53
Filesystem revision #: 1 (dynamic)
Filesystem features: has_journal ext_attr resize_inode dir_index filetype needs_recovery
extent 64bit flex_bg sparse_super large_file huge_file dir_nlink extra_isize metadata_csum
Filesystem flags: signed_directory_hash
Default mount options: user_xattr acl
Filesystem state: clean
Errors behavior: Continue
Filesystem OS type: Linux
Filesystem created: Mon Oct 14 12:34:48 2019
Last mount time: Thu Oct 24 12:27:30 2019
Last write time: Thu Oct 24 12:27:30 2019
Mount count: 50
Maximum mount count: -1
Lifetime writes: 50 GB
Default directory hash: half_md4
Directory Hash Seed: 4e7gds499c3-8532e0-452356c-432890c-d0fds43e2be81ee
Journal backup: inode blocks
Checksum type: crc32c
Checksum: 0xb054235dk
```

Server Hostname

Renaming An Ubuntu Linux Host

Renaming a host on Ubuntu is simple, just need to make some very small changes to both `/etc/hosts` and `/etc/hostname`. See the comments within the files below for more information. Once these changes are made, simply reboot the host and the changes will be applied.

```
# '/etc/hosts' should contain a line similar to the below
127.0.0.1 localhost
# Change it to the following to name the host 'alvin'
127.0.0.1 alvin
```

Similarly, the `/etc/hostname` file will contain *just* the name of the host. So, if we actually wanted to name our host 'alvin', we would change its content to reflect that.

```
alvin
```

Don't forget to reboot the host to apply the changes. Also, if you are hosting any content or running applications, be sure to save your data and stop the processes if necessary in order to avoid creating issues.

Swap Allocation

Creating swap memory for your host could prevent system or services from crashing when under heavy loads. To do this, run the following commands.

Creating Swap Files

```
# To create a 512MiB swap file -
sudo dd if=/dev/zero of=/swapfile bs=1M count=512 status=progress

# To create a 1GiB swap file -
sudo dd if=/dev/zero of=/swapfile bs=1024 count=1048576
```

Enabling Swap

After creating the swap file of the desired size, in the desired directory, we'll need to set some permissions and prepare our file to be used for swap space -

```
# Set permissions
sudo chmod 600 /swapfile

# Format file to be used for swap allocation
sudo mkswap /swapfile

# Tell our system to mount this file for swap usage
sudo swapon /swapfile
```

Adding Default Swap Entry - fstab

In short, an **fstab** entry for mounting the swap partition we created above -

```
# <device> <dir> <type> <options> <dump> <fsck>
/swapfile none swap defaults 0 0
```

Add this line to your `/etc/fstab` to mount and use this partition for swap automatically on system reboots.

For more information, see [Mounting Default Filesystems](#) or [ArchWiki - Fstab](#)

Verifying Swap Configuration

To check available system swap space, run `free -h` to see output similar to the below -

```
root@host: ~# free -h
              total        used         free       shared  buff/cache   available
Mem:          983Mi       260Mi         62Mi         0.0Ki        660Mi        560Mi
Swap:         1.0Gi         15Mi         1.0Gi
```

Alternatively, we could run 'sudo swapon --show' to see the below output

```
root@host: ~# sudo swapon --show
NAME      TYPE  SIZE  USED  PRI0
/swapfile file 1024M 15.8M  -2
```

Swappiness Values

The default swappiness value is set to 60, but to check, change, or verify your system swappiness, see the commands below

```
# Check system swappiness setting
root@host: ~# cat /proc/sys/vm/swappiness
60

# Set a new swappiness value
root@host: ~# sudo sysctl vm.swappiness=10
root@host: ~# cat /proc/sys/vm/swappiness
10
```

Swappiness Persistence

Upon setting a custom swappiness value and rebooting, your custom configuration will be lost.

Edit `/etc/sysctl.conf` to contain the line below to ensure this value is kept between system reboots

```
vm.swappiness=10
```

Removing Swap Files

First, turn off the swap file -

```
root@host: ~# sudo swapoff -v /swapfile
```

Remove the swap file entry from your `/etc/fstab` if you previously created one. If present, remove the line similar to the below

```
/swapfile swap swap defaults 0 0
```

Last, delete the swap file using `rm` -

```
root@host: ~# sudo rm /swapfile
```

Synchronizing Time Using NTP

Check out [NTP-Pool](#) for a list of pools available to different regions.

Configuration

Network Time Protocol (NTP) allows us to easily synchronize our servers with the indicated NTP host. The settings stored in `/etc/systemd/timesyncd.conf` allow us to specify which NTP server we would prefer to sync with, as well as which server(s) to use should our preferred option fail for whatever reason.

```
# This file is part of systemd.
#
# systemd is free software; you can redistribute it and/or modify it
# under the terms of the GNU Lesser General Public License as published by
# the Free Software Foundation; either version 2.1 of the License, or
# (at your option) any later version.
#
# Entries in this file show the compile time defaults.
# You can change settings by editing this file.
# Defaults can be restored by simply deleting this file.
#
# See timesyncd.conf(5) for details.

[Time]
#NTP=
#FallbackNTP=ntp.ubuntu.com
#RootDistanceMaxSec=5
#PollIntervalMinSec=32
#PollIntervalMaxSec=2048
```

The configuration above is an example of the default settings, which are commented out since

these same default settings are assumed by the Ubuntu system. If you want to change them, just remove the comment and modify their values. Below, we have modified the settings to use various servers based on our preferences.

```
# This file is part of systemd.
#
# systemd is free software; you can redistribute it and/or modify it
# under the terms of the GNU Lesser General Public License as published by
# the Free Software Foundation; either version 2.1 of the License, or
# (at your option) any later version.
#
# Entries in this file show the compile time defaults.
# You can change settings by editing this file.
# Defaults can be restored by simply deleting this file.
#
# See timesyncd.conf(5) for details.

[Time]
NTP=0.north-america.pool.ntp.org 1.north-america.pool.ntp.org
FallbackNTP=ntp.ubuntu.com 0.arch.pool.ntp.org

#RootDistanceMaxSec=5
#PollIntervalMinSec=32
#PollIntervalMaxSec=2048
```

Above, we tell systemd that we would prefer to connect to NTP servers in the following order

1. 0.north-america.pool.ntp.org
2. 1.north-america.pool.ntp.org
3. ntp.ubuntu.com
4. 0.arch.pool.ntp.org

Synchronization

If your server is for any reason out of sync, this could cause various issues down the line with services or applications. To correct this, simply synchronize with your configured NTP servers by running `sudo timedatectl set-ntp true && timedatectl status` - These commands will synchronize and then print the status of your NTP connection. Be sure to verify the information is correct, and double-check by running `date` within your bash terminal.

System Sensors

Your system likely has many sensors built in for displaying useful information on internal hardware status. For example, the commands below will help in finding the path to system temperature sensors.

```
user@host ~ $: sensors -f
coretemp-isa-0000
Adapter: ISA adapter
Package id 0:  +80.6°F (high = +176.0°F, crit = +212.0°F)
Core 0:        +73.4°F (high = +176.0°F, crit = +212.0°F)
Core 1:        +73.4°F (high = +176.0°F, crit = +212.0°F)
Core 2:        +69.8°F (high = +176.0°F, crit = +212.0°F)
Core 3:        +68.0°F (high = +176.0°F, crit = +212.0°F)

acpitz-acpi-0
Adapter: ACPI interface
temp1:         +82.0°F (crit = +221.0°F)
temp2:         +85.6°F (crit = +221.0°F)

nouveau-pci-0100
Adapter: PCI adapter
GPU core:      +0.97 V (min = +0.60 V, max = +1.27 V)
fan1:          691 RPM
temp1:         +89.6°F (high = +203.0°F, hyst = +37.4°F)
                (crit = +221.0°F, hyst = +41.0°F)
                (emerg = +275.0°F, hyst = +41.0°F)
power1:        36.13 W (crit = 275.00 mW)

asus-isa-0000
Adapter: ISA adapter
cpu_fan:       0 RPM
```

We can see that the CPU and GPU temperature sensors are known to our system as `coretemp-isa-0000` and `nouveau-pci-0100`, respectively. Run the command below to list the system path to all connected temperature devices by name, and cross-check these two outputs to gather the needed

information for your sensors.

```
user@host ~ $: for i in /sys/class/hwmon/hwmon*/temp*_input; do echo "$(<$(dirname $i)/name):  
$(cat ${i%_*}_label 2>/dev/n  
ull || echo $(basename ${i%_*})) $(readlink -f $i)"; done  
  
acpitz: temp1 /sys/devices/virtual/thermal/thermal_zone0/hwmon0/temp1_input  
acpitz: temp2 /sys/devices/virtual/thermal/thermal_zone0/hwmon0/temp2_input  
coretemp: Package id 0 /sys/devices/platform/coretemp.0/hwmon/hwmon2/temp1_input  
coretemp: Core 0 /sys/devices/platform/coretemp.0/hwmon/hwmon2/temp2_input  
coretemp: Core 1 /sys/devices/platform/coretemp.0/hwmon/hwmon2/temp3_input  
coretemp: Core 2 /sys/devices/platform/coretemp.0/hwmon/hwmon2/temp4_input  
coretemp: Core 3 /sys/devices/platform/coretemp.0/hwmon/hwmon2/temp5_input  
nouveau: temp1 /sys/devices/pci0000:00/0000:00:01.0/0000:01:00.0/hwmon/hwmon3/temp1_input
```

Displays

When attempting to manage displays, whether its the orientation or enabling / disabling, look to the man pages for `xrandr`. See the commands below for some examples.

```
# Output information on displays  
xrandr  
  
# List the output names for displays  
xrandr --output  
  
# Move DP-2 to the right of HDMI-1  
xrandr --output DP-2 --right-of HDMI-1  
``  
  
#### Timezone  
  
To see date / time, run `date`  
  
To adjust local TZ settings, run `tzselect`. Pay attention to the final output of this tool  
as it will explain how to make your change permanent. For me, I had to add the following to  
the end of my `~/.profile` :  
  
```bash
```

```
TZ='America/New_York'; export TZ
```

## Memory

Some useful commands to find information on memory usage -

```
Output various memory details
cat /proc/meminfo
#Can be used with grep, awk, etc for more specific output..
ex) Show MiB of memory available
grep -w MemAvailable: /proc/meminfo | awk '{print $2 / 1024 "MiB"}'
```

## Input Devices

Run the following to get information on input devices attached to the machine -

```
In the output shown below, my keyboard is AT Translated Set 2 keyboard
xinput list
Example output:
[Virtual core pointer id=2 [master pointer (3)]
 | ↳ Virtual core XTEST pointer id=4 [slave pointer (2)]
 | ↳ Elan Touchpad id=10 [slave pointer (2)]
[Virtual core keyboard id=3 [master keyboard (2)]
 | ↳ Virtual core XTEST keyboard id=5 [slave keyboard (3)]
 | ↳ Power Button id=6 [slave keyboard (3)]
 | ↳ Power Button id=7 [slave keyboard (3)]
 | ↳ Sleep Button id=8 [slave keyboard (3)]
 | ↳ TOSHIBA Web Camera - HD: TOSHIB id=9 [slave keyboard (3)]
 | ↳ AT Translated Set 2 keyboard id=11 [slave keyboard (3)]

Test the device..
xinput test "AT Translated Set 2 keyboard"
Example output:
key release 36
key press 40
dkey release 40
key press 50
key release 50
The output above shows me pressing / releasing keys in real time.
```

```
Exit with CTRL-C
```

## Power Supplies / AC Adapters

```
List power supplies, AC adapters -
ls -l /sys/class/power_supply/
Example output...
lrwxrwxrwx 1 root root 0 Mar 23 23:02 AC ->
../../../../devices/LNXSYSTEM:00/LNXYBUS:00/PNP0A08:00/device:00/PNP0C09:00/ACPI0003:00/power_supply/
lrwxrwxrwx 1 root root 0 Mar 23 23:02 BAT0 ->
../../../../devices/LNXSYSTEM:00/LNXYBUS:00/PNP0A08:00/device:00/PNP0C09:00/PNP0C0A:00/power_supply/B
Above, my battery is seen as BAT0, my AC port for charging is AC
```

# Unattended Upgrades

To configure linux hosts to automatically install updates and upgrades, add or edit the following lines in `/etc/apt/apt.conf.d/50unattended-upgrades`. Feel free to change the settings as you see fit.

```
Unattended-Upgrade::Mail "user@example.com";
Unattended-Upgrade::MailOnlyOnError "true";
Unattended-Upgrade::Remove-Unused-Kernel-Packages "true";
Unattended-Upgrade::Remove-Unused-Dependencies "true";
Unattended-Upgrade::Automatic-Reboot "true";
Unattended-Upgrade::Automatic-Reboot-Time "02:38";
```

At the top of `/etc/apt/apt.conf.d/50unattended-upgrades`, you'll notice the block below, be sure to follow my comments and make the changes needed

```
Unattended-Upgrade::Allowed-Origins {
 "${distro_id}:${distro_codename}";
 "${distro_id}:${distro_codename}-security";
 // Extended Security Maintenance; doesn't necessarily exist for
 // every release and this system may not have it installed, but if
 // available, the policy for updates is such that unattended-upgrades
 // should also install from here by default.
 "${distro_id}ESM:${distro_codename}";
 "${distro_id}:${distro_codename}-updates"; // <-- Uncomment this line.
 // "${distro_id}:${distro_codename}-proposed";
 // "${distro_id}:${distro_codename}-backports";
};
```

Add the following lines to `sudo vim /etc/apt/apt.conf.d/20auto-upgrades`.

```
APT::Periodic::Update-Package-Lists "1";
APT::Periodic::Unattended-Upgrade "1";
Add these two lines...
APT::Periodic::Download-Upgradeable-Packages "1";
APT::Periodic::AutocleanInterval "7";
```

Test that you can run a dry-run update using unattended-upgrades -

```
sudo unattended-upgrades --dry-run --debug
```

Also, check the logs for unattended-upgrades below

```
less /var/log/unattended-upgrades/unattended-upgrades.log
```

# Crontab

Using `crontab` to schedule tasks for the system to perform is fairly straight forward, once you get familiar with the syntax used within the configuration. Run `crontab -e` to open the file for editing, and modify to your needs using the examples below

Tell crontab where to send email alerts to by adding the following lines to any `crontab`

```
MAILTO=someuser@somedomain.com
MAILFROM=someuser@somedomain.com
```

Alternatively, to silence all emails, just provide no email with `MAILTO=' '`.

```
Schedule our system to run the test.sh script once a day -
0 0 * * * /path/to/test.sh

Syntax used for time -
* * * * * command to be executed
- - - - -
| | | | |
| | | | - - - - Day of week (0 - 7) (Sunday=0 or 7)
| | | - - - - - Month (1 - 12)
| | - - - - - Day of month (1 - 31)
| - - - - - Hour (0 - 23)
- - - - - Minute (0 - 59)

Operators used in scheduling -
(*) : This operator specifies all possible values for a field. For example, an asterisk in
the hour time field would be equivalent to every hour or an asterisk in the month field would
be equivalent to every month.
(,) : This operator specifies a list of values, for example: "1,5,10,15,20, 25".
(-) : This operator specifies a range of values, for example: "5-15" days , which is
equivalent to typing "5,6,7,8,9, ...,13,14,15" using the comma operator.
(/) : This operator specifies a step value, for example: "0-23/" can be used in the hours
field to specify command execution every other hour. Steps are also permitted after an
asterisk, so if you want to say every two hours, just use */2.
```

So, some example for various schedules -

```
To run a command once a day at midnight
0 0 * * * /path/to/unixcommand

To run /path/to/command every five minutes, every day, enter:
5 0 * * * /path/to/command

To run /path/to/command five minutes after midnight, every day, enter:
5 0 * * * /path/to/command

Run /path/to/script.sh at 2:15pm on the first of every month, enter:
15 14 1 * * /path/to/script.sh

Run /scripts/phpscript.php at 10 pm on weekdays, enter:
```

```
0 22 * * 1-5 /scripts/phpscript.php

Run /root/scripts/perl/perlscript.pl at 23 minutes after midnight, 2am, 4am ..., everyday,
enter:
23 0-23/2 * * * /root/scripts/perl/perlscript.pl

Run /path/to/unixcommand at 5 after 4 every Sunday, enter:
5 4 * * sun /path/to/unixcommand
```

Alternative, more readable but less customizable syntax for scheduling common times -

```
@reboot[]Run once, at startup.
@yearly[]Run once a year, "0 0 1 1 *".
@annually[](same as @yearly)
@monthly[]Run once a month, "0 0 1 * *".
@weekly[]Run once a week, "0 0 * * 0".
@daily[]Run once a day, "0 0 * * *".
@midnight[](same as @daily)
@hourly[]Run once an hour, "0 * * * *".
```

Useful crontab commands

```
Edit crontab configuration
crontab -e

List crontab jobs
crontab -l

Check status of cron
sudo systemctl status cron
sudo journalctl -u cron
sudo journalctl -u cron | grep backup-script.sh

Cron logs
cat /var/log/cron
tail -f /var/log/cron
grep "my-script.sh"
tail -f /var/log/cron

Backup cron
```

```
crontab -l > /nas01/backup/cron/users.root.bakup
crontab -u userName -l > /nas01/backup/cron/users.userName.bakup
```

Much of this and more information was found at [CyberCiti](#)

# Systemd Services

To define our own service with `systemd`, we need to create a `daemon.service` file. This is easily done within a few quick lines using `vim`, and should only take a few minutes.

First, we need to locate the binary for the command we want to be executed as a service. This is just good to have on-hand when defining a new service. Check where exactly your binary is using `which <command>`, seen below

```
which hexo
/home/hexouser/.nvm/versions/node/v20.9.9/bin/hexo
```

Now we know exactly where the binary that we execute is when we run the `hexo` command, and we will use it within the `hexo.service` file we create below, so be sure to have it handy.

To create a user service, place the `hexo.service` file within the `$HOME/.config/systemd/hexouser/` directory. This will allow the user to manage the service without `sudo` by running `systemd --user start name.service`

Create a service file like the one below for `hexo` by running `sudo vim /etc/systemd/system/hexo.service`. If you are defining a service for something else, just rename this file accordingly.

```
[Unit]
Description=Personal hexo blog service
After=network.target

[Service]
Type=simple
Another Type: forking
User=hexouser
WorkingDirectory=/home/hexouser/hexosite
ExecStart=/home/hexouser/.nvm/versions/node/v29.9.9/bin/hexo server --cwd
/home/hexoroot/hexosite
```

```
ExecStop=/bin/kill -TERM $MAINPID
ExecReload=/bin/kill -HUP $MAINPID
Restart=on-failure
Other restart options: always, on-abort, etc

The install section is needed to use
`systemctl enable` to start on boot
For a user service that you want to enable
and start automatically, use `default.target`
For system level services, use `multi-user.target`
[Install]
WantedBy=multi-user.target
WantedBy=graphical.target
```

When making changes to a service, you need to run `sudo systemctl daemon-reload` between edits to apply your changes before restarting your service. Once the above file is created within `/etc/systemd/system/hexo.service` we can start our hexo blog using systemd by running the usual commands

```
Start your new service
sudo systemctl start hexo.service
Enable your service to start automatically on reboot or crashing
sudo systemctl enable hexo.service
Check on your service
sudo systemctl status hexo.service
```

We can even check on our logs using `journalctl`

```
sudo journalctl -u hexo
journalctl --user-unit hexo
```