

---

# **paradrop Documentation**

***Release 0.4***

**Paradrop Labs**

February 16, 2017



<b>1</b>	<b>Getting Started</b>	<b>3</b>
1.1	Environment setup . . . . .	3
1.2	Activating your Router . . . . .	3
1.3	Installing Chutes . . . . .	4
<b>2</b>	<b>Chute Development</b>	<b>5</b>
2.1	Structure . . . . .	5
2.2	Dockerfile . . . . .	5
2.3	Persistent Data . . . . .	6
2.4	System Information . . . . .	6
<b>3</b>	<b>Developing Light Chutes</b>	<b>7</b>
3.1	Structure . . . . .	7
3.2	paradrop.yaml . . . . .	7
3.3	Persistent Data . . . . .	8
<b>4</b>	<b>Chute Deployment</b>	<b>9</b>
<b>5</b>	<b>Installing on the Intel NUC</b>	<b>11</b>
5.1	Hardware and software requirements . . . . .	11
5.2	Preparing for installation . . . . .	11
5.3	Boot from the Live USB flash drive . . . . .	11
5.4	Flash ParaDrop . . . . .	12
5.5	First boot . . . . .	12
5.6	Activate the Router . . . . .	12
<b>6</b>	<b>The Paradrop Instance System</b>	<b>15</b>
6.1	Build System . . . . .	15
6.2	Snappy Confinement . . . . .	16
6.3	Documentation and tests . . . . .	17
<b>7</b>	<b>Architecture</b>	<b>19</b>
<b>8</b>	<b>Known Issues</b>	<b>21</b>
8.1	Issues using pdbuild.sh . . . . .	21
8.2	Issues with the hardware or operating system . . . . .	22
<b>9</b>	<b>Frequently Asked Questions</b>	<b>23</b>

<b>10</b>	<b>Paradrop</b>	<b>25</b>
<b>11</b>	<b>The Paradrop workflow</b>	<b>27</b>
<b>12</b>	<b>Getting Started</b>	<b>29</b>
<b>13</b>	<b>Where to go from here?</b>	<b>31</b>
<b>14</b>	<b>What if I don't have Ubuntu?</b>	<b>33</b>

This section of the documentation is devoted to describing the apps that run on Paradrop.

If this is your first time seeing Paradrop, please start with the [Getting Started](#) page.

Contents:



---

## Getting Started

---

This will quickly take you through the process of bringing up a Hello World chute in a virtual machine on your computer.

*NOTE:* These instructions assume you are running Ubuntu. The steps to launch a virtual machine may be different for other environments.

### Environment setup

These steps will download our router image and launch it in a virtual machine.

1. Install required packages:

```
sudo apt-get install qemu-kvm
```

2. Download the latest image (paradrop\_router.img.gz) from [our releases](#).

3. Extract the image:

```
gunzip paradrop_router.img.gz
```

4. Launch the VM:

```
sudo kvm -m 1024 \
-netdev user,id=net0,hostfwd=tcp::8000-:8000,hostfwd=tcp::80-:80 \
-device virtio-net-pci,netdev=net0 -drive file=paradrop_router.img,format=raw
```

Please note: there is no username/password to log into the system. Please follow the steps in the next section to access your router through [paradrop.org](#).

### Activating your Router

Follow these steps the first time you start a new physical or virtual Paradrop router. Activation associates the router with your account on [paradrop.org](#) so that you can manage the router and install chutes from the chute store.

1. Make an account on [paradrop.org](#) if you do not have one.
2. From the Routers tab, click Create Router. Give your router a unique name and an optional description to help you remember it and click Submit.
3. On the router page, find the “Router ID” and “Password” fields. You will need to copy this information to the router so that it can connect.

4. Open the router portal at <http://localhost>. Click the “Activate the router witha ParaDrop Router ID” button and enter the information from the [paradrop.org](http://paradrop.org) router page. If the activation was successful, you should see checkmarks appear on the “WAMP Router” and “ParaDrop Server” lines. You may need to refresh the page to see the update.

## Installing Chutes

1. Make an account on [paradrop.org](http://paradrop.org) and make sure you have an activated, online router.
2. Go to the Chute Store tab on [paradrop.org](http://paradrop.org). There you will find some public chutes such as the hello-world chute. You can also create your own chutes here.
3. Click on the hello-world chute, click Install, click your router’s name to select it, and finally, click Install.
4. That will take you to the router page again where you can click the update item to monitor its progress. When the installation is complete, an entry will appear under the Chutes list.
5. The hello-world chute starts a webserver, which is accessible at <http://localhost:8000>. Once the installation is complete, test it in a web browser.



---

## Chute Development

---

Minimally, a chute has a Dockerfile, which contains instructions for building and preparing the application to run on Paradrop. A chute will usually also require scripts, binaries, configuration files, and other assets. For integration with the Paradrop toolset, we highly recommend developing a chute as a [GitHub](#) project, but other organization methods are possible.

We will examine the [hello-world](#) chute as an example of a complete Paradrop application.

### Structure

Our hello-world chute is a git project with the following files:

```
chute/index.html
Dockerfile
README.md
```

The top-level contains a README and a special file called “Dockerfile”, which will be discussed below. As a convention, we place files that will be used by the running application in a subdirectory called “chute”. This is not necessary but helps keep the project organized. Valid alternatives include “src” or “app”.

### Dockerfile

The Dockerfile contains instructions for building and preparing an application to run on Paradrop. Here is a minimal Dockerfile for our hello-world chute:

```
FROM nginx
ADD chute/index.html /usr/share/nginx/html/index.html
```

#### FROM nginx

The FROM instruction specifies a base image for the chute. This could be a Linux distribution such as “ubuntu:14.04” or an standalone application such as “nginx”. The image name must match an image in the Docker public registry. We recommend choosing from the [official repositories](#). Here we use “nginx” for a light-weight web server.

#### ADD chute/index.html index.html

The ADD instruction copies a file or directory from the source repository to the chute filesystem. This is useful for installing scripts or other files required by the chute and are part of the source repository. The <source> path should be inside the repository, and the <destination> path should be an absolute path or a path inside the chute’s working directory. Here we install the index.html file from our source repository to the search directory used by nginx.

Other useful commands for building chutes are RUN and CMD. For a complete reference, please visit the official [Dockerfile reference](#).

Here is an alternative implementation of the hello-world Dockerfile that demonstrates some of the other useful instructions.

```
FROM ubuntu:14.04
RUN apt-get update && apt-get install -y nginx
ADD chute/index.html /usr/share/nginx/html/index.html
EXPOSE 80
CMD ["nginx", "-g", "daemon off;"]
```

Here we use a RUN instruction to install nginx and a CMD instruction to set nginx as the command to run inside the chute container. Using “ubuntu:14.04” as the base image gives access to any packages that can be installed through apt-get.

## Persistent Data

Each running chute has a persistent data storage that is not visible to other chutes. By default the persistent data directory is named “data” inside the chute’s filesystem. Files stored in this directory will remain when upgrading or downgrading the chute and are only removed when uninstalling the chute.

## System Information

The Paradrop instance tools share system information with chutes through a read-only directory named “/paradrop”. Chutes that are configured with a WiFi access point will find a file in this directory that lists wireless clients. In future versions there will also be information about Bluetooth and other wireless devices.

### dnsmasq-wifi.leases

This file lists client devices that have connected to the WiFi AP and received a DHCP lease. This is a plain text file with one line for each device containing the following space-separated fields.

1. DHCP lease expiration time (seconds since Unix epoch).
2. MAC address.
3. IP address.
4. Host name, if known.
5. Client ID, if known; the format of this field varies between devices.

The following example shows two devices connected to the chute’s WiFi network.

```
1480650200 00:11:22:33:44:55 192.168.128.130 android-ffeeddccbbaa9988 *
1480640500 00:22:44:66:88:aa 192.168.128.170 someones-iPod 01:00:22:44:66:88:aa
```

---

## Developing Light Chutes

---

Light chutes build and install the same way as normal chutes and can do many of the same things. However, they make use of prebuilt base images that are optimized for different programming languages.

Light chutes offer these advantages over normal, *heavy* chutes.

- **Safety:** Light chutes have stronger confinement properties, so you can feel safer installing a light chute written by a third party developer.
- **Fast installation:** Light chutes use a common base image that may already be cached on the router, so installation can be very fast.
- **Simplicity:** You do not need to learn how to write and debug a Dockerfile to develop a chute. Instead, you can use the package management tools you may already be using (e.g. `package.json` for npm and `requirements.txt` for pip).
- **Portability:** With ARM support coming soon for ParaDrop, your light chutes will most likely run on ARM with extra work on your part. This is not the case for normal chutes that use a custom Dockerfile.

We will look at the `node-hello-world` chute as an example of a light chute for ParaDrop.

### Structure

Our hello-world chute is a git project with the following files:

```
README.md
index.js
package.json
paradrop.yaml
```

The project contains the typical files for a node.js project as well as a special file called “paradrop.yaml”.

### paradrop.yaml

The `paradrop.yaml` file contains information that ParaDrop needs in order to run the chute. Here are the contents for the hello-world example:

```
version: 1
type: light
use: node
command: node index.js
```

All of these fields are required and very simple to use.

**version: 1**

This specifies the version of the `paradrop.yaml` schema in order to allow future changes without breaking existing chutes. You should specify version *1*.

**type: light**

This indicates that we are building a *light* chute.

**use: node**

This indicates that we are using the *node* base image for this chute. You should choose the base image appropriate for your project. Supported images are: *node* and *python2*.

**command: node index.js**

This line indicates the command for starting your application. You can either specified it this way, as a string with spaces between the parameters, or you can use a list of strings. The latter format would be particularly useful if your parameters include spaces. Here is an example:

```
command:
- node
- index.js
```

## Persistent Data

Each running chute has a persistent data storage that is not visible to other chutes. By default the persistent data directory is named “/data” inside the chute’s filesystem. Files stored in this directory will remain when upgrading or downgrading the chute and are only removed when uninstalling the chute.

---

# Chute Deployment

---

Starting with version 0.2, developers should use our website, [paradrop.org](https://paradrop.org) for creating chutes and installing them on Paradrop routers. The command line tools from Paradrop 0.1 are not currently supported.

As our website is under development, please refer to the instructions on the website or contact us for help with making and installing a chute.



---

## Installing on the Intel NUC

---

These instructions will help you install the ParaDrop daemon on the Intel NUC platform. At the end of this process, you will have a system ready for installing chutes.

We have specifically tested this process on the Skull Canyon (NUC6i7KYK) platform, which we recommend for high performance edge-computing needs.

### Hardware and software requirements

- **Intel NUC Skull Canyon NUC6i7KYK**
  - The Intel NUC devices generally do not come with memory or storage pre-installed.
  - Memory: we recommend at least one 8 GB DDR4 SODIMM.
  - Storage: we have generally found one 16 GB SD card to be sufficient for our storage needs, but we recommend using one MX300 M.2 SSD card for the higher read and write speeds.
  - We recommend updating the BIOS on the NUC. Follow the instructions on [the Intel support site](#).
- 2 USB 2.0 or 3.0 flash drives (each 4 GB minimum)
- A monitor with an HDMI interface
- A network connection with Internet access
- An [Ubuntu Desktop 16.04.1 LTS](#) image.
- A [ParaDrop](#) disk image.

### Preparing for installation

1. Download the Ubuntu Desktop image and prepare a bootable USB flash drive.
2. Download the ParaDrop disk image and copy the file to the second flash drive.

### Boot from the Live USB flash drive

1. Insert the Live USB Ubuntu Desktop flash drive in the NUC.
2. Start the NUC and push F10 to enter the boot menu.

3. Select the USB flash drive as a boot option.
4. Select “Try Ubuntu without installing”.

## Flash ParaDrop

1. Once the system is ready, insert the second USB flash drive which contains the ParaDrop disk image.
2. Open a terminal and run the following command, where <disk label> is the name of the second USB flash drive. You may wish to double-check that /dev/sda is the desired destination **before running dd**.

```
zcat /media/ubuntu/<disk label>/paradrop_router.img.gz | sudo dd of=/dev/sda bs=32M status=progress
```

3. Reboot the system and remove all USB flash drives when prompted to do so.

## First boot

1. At the Grub menu, press ‘e’ to edit the boot options.
2. Find the line that begins with “linux” and append the option “nomodeset”. It should look like “linux (loop)/kernel.img \$cmdline nomodeset”. Adding this option will temporarily fix a graphics issue that is known to occur with the Intel NUC.
3. Press F10 to continue booting.
4. Press Enter when prompted, and follow the instructions on the screen to configure Ubuntu Core. If you have an Ubuntu One account. By connecting your Ubuntu One account, you will be able to login via SSH with the key(s) attached to your account. Otherwise, if you do not have an Ubuntu One account or do not wish to use it, you may enter “[info@paradrop.io](mailto:info@paradrop.io)” as your email address. You will still be able to manage your router and install chutes through [paradrop.org](http://paradrop.org) either way.
5. Take note of the IP address displayed on the screen. You will need this address for the next step, activating the router. For example, the message below indicates that the router has IP address 10.42.0.162.

```
Congratulations! This device is now registered to info@paradrop.io.

The next step is to log into the device via ssh:

ssh paradrop@10.42.0.162
...
```

## Activate the Router

Follow these steps the first time you start a new physical or virtual ParaDrop router. Activation associates the router with your account on [paradrop.org](http://paradrop.org) so that you can manage the router and install chutes from the chute store.

1. Make an account on [paradrop.org](http://paradrop.org) if you do not have one.
2. From the Routers tab, click Create Router. Give your router a unique name and an optional description to help you remember it and click Submit.
3. On the router page, find the “Router ID” and “Password” fields. You will need to copy this information to the router so that it can connect.



4. Open the router portal at [http://<router\\_address>](http://<router_address>). Click the “Activate the router with a ParaDrop Router ID” button and enter the information from the [paradrop.org](http://paradrop.org) router page. If the activation was successful, you should see checkmarks appear on the “WAMP Router” and “ParaDrop Server” lines. You may need to refresh the page to see the update.



---

## The Paradrop Instance System

---

This section focuses on the *Instance Tools*. This is the set of daemons and tools required to allow the Paradrop platform to function on virtual machines and real hardware. Use the information below to learn about Snappy Ubuntu and how we leverage it to create next generation smart routers.

Contents:

### Build System

Paradrop includes a set of build tools to make development as easy as possible.

Currently this system takes the form of a bash script that automates installation and execution, but in time this may evolve into a published python package. This page outlines the steps required to manually build the components required to develop with paradrop.

Components in the build process:

- *Installing and running Ubuntu Snappy*
- *Building paradrop*
- *Installing paradrop*
- **‘Creating chutes’\_**

We recommend using Ubuntu 14.04 as the build environment for this version of Paradrop. Ubuntu 16.04 will not work because the snappy development tools have changed. The next release of Paradrop will use the new tools.

You will only need to follow these instructions if you will be making changes to the Paradrop instance tools. Otherwise, you can use our pre-built Paradrop disk image. Please visit the [Getting Started](#) page.

### Installing and running Ubuntu Snappy

**Snappy** is an Ubuntu release focusing on low-overhead for a large set of platforms. These instructions are for getting a Snappy instance up and running using ‘kvm’.

Download and unzip a snappy image:

```
wget http://releases.ubuntu.com/15.04/ubuntu-15.04-snappy-amd64-generic.img.xz
unxz ubuntu-15.04-snappy-amd64-generic.img.xz
```

Launch the snappy image using kvm:

```
kvm -m 512 -redir :8090::80 -redir :8022::22 ubuntu-15.04-snappy-amd64-generic.img
```

Connect to local instance using ssh:

```
ssh -p 8022 ubuntu@localhost
```

## Building paradrop

Snappy is a closed system (by design!). Arbitrary program installation is not allowed, so to allow paradrop access to the wide world of `pypi` the build system relies on two tools.

- `virtualenv` is a tool that creates encapsulated environments in which python packages can be installed.
- `pex` can compress python packages into a zip file that can be executed by any python interpreter.
- `snappy` is a tool for building snap packages. *Note:* Ubuntu 16.04 uses `snappy` instead, which produces incompatible packages.

First, set `DEV_MACHINE_IP=paradrop.org` in `pdbuild.conf`. The build script will refuse to run if this variable is not set.

Install the necessary development tools:

```
./pdbuild.sh setup
```

Build the Paradrop snap package:

```
./pdbuild.sh build
```

## Installing paradrop

Install dependencies on the virtual machine:

```
./pdbuild.sh install_deps
```

Install the newly created Paradrop snap package:

```
./pdbuild.sh install_dev
```

## Snappy Confinement

Snappy confines running applications in two ways: directory isolation and mandatory access control (MAC). Directory isolation means the application cannot leave its installed directory. MAC means the application cannot execute any system commands or access any files it does not have explicit, predetermined permissions to.

MAC is the more serious hurdle for paradrop development. Snaps declare permissions through an [AppArmor profile](#).

## Getting started with Profile Generation

Install tools and profiles:

```
sudo apt-get install apparmor-profiles apparmor-utils
```

List active profiles:

```
sudo apparmor_status
```

Profiles in complain mode log behavior, while those in enforce mode actively restrict it.

The following steps assume paradrop is installed on the system and not on a virtualenv.

Create a new, blank profile:

```
cd /etc/apparmor.d/
sudo aa-autodep paradrop
```

Use aa-complain to put the profile in complain mode:

```
sudo aa-complain paradrop
```

Exercise the application! AppArmor will surreptitiously watch the program in the background and log all behavior. Once finished, use the following command to go through the resulting requests, approve or deny them, and autogenerate a profile:

```
sudo aa-logprof
```

## Documentation and tests

Documentation is handled by [sphinx](#) and [readthedocs](#).

Testing is a joint effort between [nosetests](#), [travis-ci](#), and [coveralls](#).

### Documentation

Information about docs creation, management, and display.

Sphinx reads files in [reStructuredText](#) and builds a set of HTML pages. Every time a new commit is pushed to github, [readthedocs](#) automatically updates documentation.

Additionally, sphinx knows all about python! The directives `automodule`, `autoclass`, `autofunction` and more instruct sphinx to inspect the code located in `src/` and build documentation from the docstrings within.

For example, the directive `.. automodule:: paradrop.backend` will build all the documentation for the given package. See [google](#) for more instructions.

All docstring documentation is rebuilt on every commit (unless there's a bug in the code.) Sphinx does not, however, know about structural changes in code! To alert sphinx of these changes, use the `autodoc` feature:

```
sphinx-apidoc -f -o docs/api paradrop/paradrop/
```

This scans packages in the `src/paradrop` directory and creates `.rst` files in `docs/api`. The root file `index.rst` links to `modules.rst`, connecting the newly generated api code with the main documentation.

To create the documentation locally, run:

```
cd docs
make html
python -m SimpleHTTPServer 9999
```

Open your web browser of choice and point it to [http://localhost:9999/\\_build/html/index.html](http://localhost:9999/_build/html/index.html).

## Testing

As mentioned above, all testing is automatically run by travis-ci, a continuous integration service.

To manually run tests, install nosetest:

```
pip install nose
```

Install the required packages:

```
pip install -r docs/requirements.txt
```

Run all tests:

```
nosetests
```

Well thats easy. How does nose detect tests? All tests live in the `tests/` directory. Nose adheres to a simple principle: anything marked with `test` in its name is most likely a test. When writing tests, make sure all functions begin with `test`.

Coverage analysis detects how much of the code is used by a test suite. If the result of the coverage is less than 100%, someone slacked. Install coveralls:

```
pip install coveralls
```

Run tests with coverage analysis:

```
nosetests --with-coverage --cover-package=paradrop
```

---

## Architecture

---

This section details some of the non-obvious architectural features of paradrop.

This is a work in progress. Check back later!





---

## Known Issues

---

Please check here for issues during setup.

### Issues using pdbuild.sh

These issues are related to the *Instance Tools* found on github.

#### Issue 1: pdbuild.sh install fails

Docker snap is missing inside of virtual router, run `pdbuild.sh install_deps`:

```
issues while running ssh command: Installing /tmp/paradrop_0.1.0_all.snap
2015/08/11 21:22:57 Signature check failed, but installing anyway as requested
/tmp/paradrop_0.1.0_all.snap failed to install: missing frameworks: docker
```

#### Issue 2: pdbuild.sh install fails

This is a known issue for the Paradrop team, if you get this please email us at [developers@paradrop.io](mailto:developers@paradrop.io):

```
Installing paradrop_0.1.0_all.snap from local environment

issues while running ssh command: Installing /tmp/paradrop_0.1.0_all.snap
2015/08/11 21:29:48 Signature check failed, but installing anyway as requested
/tmp/paradrop_0.1.0_all.snap failed to install: [start paradrop_pdconfd_0.1.0.service]
failed with exit status 1: Job for paradrop_pdconfd_0.1.0.service failed.
See "systemctl status paradrop_pdconfd_0.1.0.service" and "journalctl -xe" for details.
```

#### Issue 3: pdbuild.sh up fails

This is very common and will happen if you delete your VM and setup a fresh one, which will have a different key. The solution is simple and is stated in the error message. Follow the instructions to remove the key from your `known_hosts` file.

```
Failed to setup keys: failed to setup keys: issues while running ssh command:
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@      WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!      @
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
```

```
IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY!  
Someone could be eavesdropping on you right now (man-in-the-middle attack)!  
It is also possible that a host key has just been changed.  
The fingerprint for the ECDSA key sent by the remote host is  
e6:ec:b1:93:7d:91:84:50:19:36:14:8e:ce:ef:6a:0b.  
Please contact your system administrator.
```

## Issues with the hardware or operating system

### Issue 1: Docker fails to start after a reboot

This can happen if the ‘docker.pid’ file was not properly cleaned up, which causes the docker daemon to conclude that it is already running.

To fix this, remove the pid file on the router and reboot.

```
sudo rm /var/lib/apps/docker/1.11.2/run/docker.pid  
sudo reboot
```

### Issue 2: WiFi devices are not detected after a reboot

Occasionally, when routers start up the WiFi devices are not detected properly. When this happens the command `iw dev` will display nothing instead of the expected devices. This is usually remedied by rebooting.

---

## Frequently Asked Questions

---

Please check here for any FAQ's.



---

# Paradrop

---

Paradrop is a software platform that enables apps to run on Wi-Fi routers. We call these apps “chutes” like a parachute. The name Paradrop comes from the fact that we are enabling the ability to “drop” supplies and resources (“apps”) into a difficult and not well-travelled environment - the home.

Paradrop runs on top of [Snappy Ubuntu](#), a trimmed-down and secured operating system that can run on ARM and x86. We also enable our apps through containerization by leveraging [Docker](#).



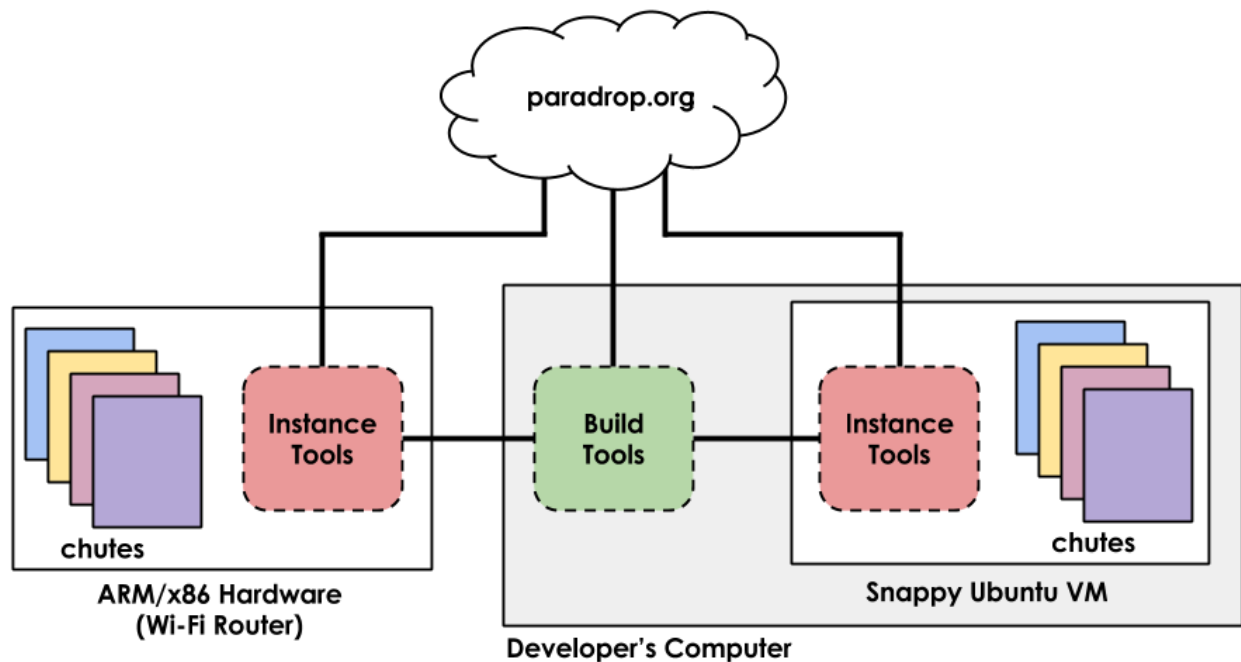
---

## The Paradrop workflow

---

There are two components to the Paradrop platform:

- The **build tools** - our CLI that enables registration, login, and control. With version 0.2 and up, Paradrop routers can be managed through our **cloud management** service instead of the CLI.
- The **instance tools** - our configuration daemons and tools to launch apps on hardware.



As you can see from the image above, we will refer to *Build Tools* when we talk about the CLI program running on your development computer that controls and communicates with the rest of the Paradrop platform. Treat this tool as your window into the rest of the Paradrop world. Our *Instance Tools* leverage programs like Docker to allow Paradrop apps to run on router hardware. This “hardware” could be a Raspberry Pi, or even a virtual machine on your computer that acts as a router (which is why we call it an Instance sometimes). Using Paradrop, you can actually plug in a USB Wi-Fi adapter and turn a virtual machine on your computer into a real router with our platform!





---

### Getting Started

---

Please visit the [Getting Started](#) page for a quick introduction to Paradrop.



---

### Where to go from here?

---

We have advanced app examples found under Apps on Paradrop. If you are interested in working on the instance side of paradrop (our github code) than check out: [The Paradrop Instance System](#).



---

## What if I don't have Ubuntu?

---

With version 0.2 and up, all of Paradrop's capabilities can be managed through our web-based service at [paradrop.org](https://paradrop.org).