
stackd.io Documentation

Release 0.7.8.dev20160708213232

Clark Perkins

July 08, 2016

1	Installation Overview	3
2	Manual Install	5
3	Using the Amazon AMI	13
4	LDAP Guide	15
5	Webserver Guide	17
6	Contact Information	21
7	Contributor Guidelines	23

stackd.io is a modern cloud deployment and provisioning framework for everyone. Its purpose is to provide a common platform for deploying and configuring hardware on **any** cloud platform. We currently only support AWS EC2, but the driver framework is expandable to support other cloud providers.

stackd.io is a [Django](#) project, and uses [Salt](#) for its back-end configuration management.

Installation Overview

There are two main options for installation:

1.1 Amazon AMI

Reading through long install guides and executing each and every command can be time consuming and error prone. If you would rather just run a script to do a lot of this for you, we have a script to build an AMI for you. Keep in mind that the script is somewhat opinionated and won't let you make many decisions (you're free to modify it to suit your needs though!) Here's a list of things it will do:

- Install all of the necessary stuff (MySQL, python, virtualenv, tons of packages, etc)
- Create a `stackdio` virtualenv at `/usr/share/stackdio`
- Install `stackdio` and its python dependencies
- Install and configure Nginx
- Install and configure `supervisord` to run `gunicorn`, `celery`, and `salt-master`
- Create an `admin` `stackdio` user

Check it out here: [Using the Amazon AMI](#)

1.2 Manual Install

If you'd rather have a more custom-fitted installation that fits your needs, check out [Manual Install](#) instead.

Manual Install

This guide is intended to quickly march you through the steps of installing and running stackd.io and its dependencies. We're not intending to be complete or provide you with everything needed for a production-ready install, we may make some assumptions you don't agree with, and there may be things we missed. If you feel anything is out of the ordinary, a bit confusing, or just plain missing, please [contact us](#).

2.1 1. A database

stackd.io needs a relational database to store internal information. Since it's built on Django, it inherently supports many different database servers. Mysql will do fine, but if you'd prefer something else, it is beyond the scope of this guide to install it. For more information on Django's database support, see: <https://docs.djangoproject.com/en/1.8/ref/databases/>

The OS-specific prep of your choice (below) will walk you through installing mysql.

2.2 2. OS-specific preparation

Warning: You must follow the steps in one of the following prep guides for the OS you're installing stackd.io in.

Follow one of the individual guides below to prepare your particular environment for stackd.io. Once you finish, come back here and continue on.

2.2.1 Preparing Ubuntu for stackd.io installation

The steps below were written using Ubuntu 13.10 from a Ubuntu-provided AMI on Amazon Web Services (AWS). The exact AMI we used is `ami-2f252646`, and you should be able to easily launch an EC2 instance using this AMI from the [AWS EC2 Console](#).

Prerequisites

All of these steps require `root` or `sudo` access. Before installing anything with `apt-get` you should run `apt-get update` first.

MySQL

Note: Please skip this section if you are using a different database or already have a supported database server running elsewhere.

Install MySQL server:

```
sudo apt-get install mysql-server mysql-client

# When prompted, provide a password for the root user to access the MySQL server.
```

Below we'll create a `stackdio` database and grant permissions to the `stackdio` user for that database.

WARNING: we're not focusing on security here, so the default MySQL setup definitely needs to be tweaked, passwords changed, etc., but for a quick-start guide this is out of scope. Please, don't run this as-is in production :)

```
echo "create database stackdio; \
grant all on stackdio.* to stackdio@'localhost' identified by 'password';" | \
mysql -hlocalhost -uroot -ppassword
```

virtualenvwrapper

```
# install the package
sudo apt-get install virtualenvwrapper

# post-install step for virtualenvwrapper shortcuts
source /etc/bash_completion.d/virtualenvwrapper
```

Core requirements

- gcc and other development tools
- git
- mysql-devel
- swig
- python-devel
- rabbitmq-server

To quickly get up and running, you can run the following to install the required packages.

```
# Install requirements needed to install stackd.io
sudo apt-get install python-dev libssl-dev libncurses5-dev libyaml-dev swig nodejs npm \
libmysqlclient-dev rabbitmq-server git nginx libldap2-dev libsasl2-dev
```

```
# Link nodejs over to node - bower will complain otherwise
sudo ln -s /usr/bin/nodejs /usr/bin/node
```

Next Steps

You're now finished with the Ubuntu-specific requirements for `stackd.io`. You can head back over to the [Manual Install](#) and continue the installation of `stackd.io`.

2.2.2 Preparing CentOS for stackd.io installation

The steps below were written using CentOS 6.4 from a CentOS-provided AMI on Amazon Web Services (AWS). The exact AMI we used is `ami-bf5021d6`, and you should be able to easily launch an EC2 instance using this AMI from the [AWS Marketplace](#).

Prerequisites

All of the CentOS-provided AMIs have SELinux and iptables enabled. We disabled both of these to be as straight forward as possible during this guide. SELinux causes issues that are beyond the scope of the guide, and we disabled iptables because we leverage EC2's security groups for firewall access.

iptables

Let's just turn it off for now. Please note, if you're using EC2 or some other cloud provider that has firewall rules enabled by default, you will need to configure the particular firewall rules to gain access to the web server we'll start in the guide. The default port for the webserver is 8000, so open this port up at a minimum. (Port 22 for SSH will obviously be needed as well.)

```
sudo service iptables stop
```

SELinux

Getting things working using SELinux could be an entirely separate guide. For our purposes, it's completely out of scope, so we're going to disable it.

Note: You will be required to restart the machine during this step.

```
# Edit /etc/sysconfig/selinux and make sure the line beginning
# with SELINUX looks like:
SELINUX=disabled

# If it was already disabled you can skip the following, however
# if you switched the policy from anything other than 'disabled'
# you need to relabel the filesystem to remove the garbage that
# SELinux has added. This *requires* a restart to take effect.
touch /.autorelabel
reboot

# When the machine is back up, you can confirm SELinux is not
# running
>>> selinuxenabled
>>> echo $?
>>> 1

# If the output is 1 you're good to go.
```

EPEL

```
sudo rpm -Uvh http://mirror.steadfast.net/epel/6/i386/epel-release-6-8.noarch.rpm
```

MySQL

Note: Please skip this section if you are using a different database or already have a supported database server running elsewhere.

Install MySQL server:

```
sudo yum install mysql-server
```

Start MySQL server:

```
sudo service mysqld start
```

Below we'll create a `stackdio` database and grant permissions to the `stackdio` user for that database.

Warning: We're not focusing on security here, so the default MySQL setup definitely needs to be tweaked, passwords changed, etc., but for a quick-start guide this is out of scope. Please, don't run this as-is in production :)

```
echo "create database stackdio; \  
grant all on stackdio.* to stackdio@'localhost' identified by 'password';" | \  
mysql -h localhost -u root
```

virtualenvwrapper

```
# install the package  
sudo yum install python-virtualenvwrapper  
  
# Update the user's ~/.bash_profile to enable virtualenvwrapper  
# You're using the stackdio user, right? :)  
echo "source /usr/bin/virtualenvwrapper.sh" >> ~/.bash_profile  
  
# re-source the .bash_profile  
. ~/.bash_profile
```

Core requirements

- gcc and other development tools
- git
- mysql-devel
- swig
- python-devel
- rabbitmq-server
- nginx

To quickly get up and running, you can run the following to install the required packages.

```
# Install the development tools group
sudo yum groupinstall "Development Tools"

# Install the other requirements needed to install stackd.io
sudo yum install git mysql-devel swig python-devel rabbitmq-server nginx nodejs npm
```

Next Steps

You're now finished with the CentOS-specific requirements for stackd.io. You can head back over to the [Manual Install](#) and continue the installation of stackd.io.

2.3 3. Create a virtualenv

Let's create a virtualenv to install stackd.io into:

```
mkvirtualenv stackdio
```

The virtualenv should automatically activate when you create it. If you exit your current shell and come back later to work on stackdio and find things not working as expected you probably need to activate the virtualenv again. To do this, `virtualenvwrapper` gives you the `workon` command:

```
workon stackdio
```

2.4 4. Install bower

In your terminal, run the following command to install bower:

Note: You must have previously installed npm/node from the OS specific preparation

```
sudo npm install -g bower
```

2.5 5. Install stackd.io

Note: Double-check that your virtualenv is activated or else this will probably complain that you don't have permissions to install (because it's trying to install into the global python site-packages directory which we don't want!)

There's two options for installing here. We recommend pulling the latest version from [pypi](#) with `pip`, like this:

```
workon stackdio # Activate the virtualenv
pip install stackdio-server[production,mysql]
```

2.6 6. Configuration

After the install, you'll have a `stackdio` command available to interact with much of the platform. First off, we need to configure stackd.io a bit. The `stackdio init` command will prompt you for several pieces of information.

If you followed all steps above verbatim, then all defaults may be accepted, but if you deviated from the path you will need to provide the following information:

- an existing user on the system that will run everything (it will default to the `stackdio` user)
- an existing location where stackd.io can store its data (the default is `$HOME/.stackdio/storage` and will be created for you if permissions allow)
- a database DSN that points to a running database you have access to (if you're using the MySQL install from above, the default `mysql://stackdio:password@localhost:3306/stackdio` is appropriate)

```
stackdio init
```

Now, let's populate are database with a schema:

```
stackdio manage.py migrate
```

2.7 7. stackd.io users

2.7.1 LDAP

stackd.io can easily integrate with an LDAP server. See our [LDAP Guide](#) for more information on configuring stackd.io to work with LDAP. If you choose to go the LDAP route, you can skip this entire section because users who successfully authenticate and are members of the right groups via LDAP will automatically be created in stackd.io.

2.7.2 Non-LDAP admin user

Admin users in stackd.io have less restriction to various pieces of the platform. For example, only admin users are allowed to create and modify cloud providers and profiles that other users can use to spin up their stacks.

Note: You will need at least one admin user to configure some key areas of the system.

```
stackdio manage.py createsuperuser
```

```
# and follow prompts...
```

2.7.3 Non-LDAP regular users

When not using LDAP, the easiest way to create new non-admin users is to use the built-in Django admin interface. First we need the server to be up and running so keep following the steps below and we'll come back to adding users later.

2.8 8. Web server configuration

For this guide, we'll use the `stackdio` command to generate the necessary configuration for Nginx to serve our static content as well as proxying the Python app through gunicorn.

To configure Nginx for CentOS:

```
# CENTOS

# add execute permissions to the user's home directory for static content to serve correctly
chmod +x ~/

stackdio config nginx | sudo tee /etc/nginx/conf.d/stackdio.conf > /dev/null

# rename the default server configuration
sudo mv /etc/nginx/conf.d/default.conf /etc/nginx/conf.d/default.conf.bak
```

To configure Nginx for Ubuntu:

```
# UBUNTU ONLY
stackdio config nginx | sudo tee /etc/nginx/sites-available/stackdio > /dev/null
sudo ln -s /etc/nginx/sites-available/stackdio /etc/nginx/sites-enabled

# remove the default configuration symlink
sudo rm /etc/nginx/sites-enabled/default
```

After this, generate the static content we'll need to serve:

```
stackdio manage.py collectstatic --noinput
```

and finally, start Nginx:

```
sudo service nginx restart
```

2.9 9. RabbitMQ, celery, and salt

Start the rabbitmq server:

```
sudo service rabbitmq-server start
```

For celery and salt-master, we'll be using supervisord. The required packages should already be installed, so we'll just need to configure supervisor and start the services.

```
# generate supervisord configuration that controls gunicorn, celery, and salt-master and store it in
stackdio config supervisord > ~/.stackdio/supervisord.conf

# launch supervisord and start the services
supervisord -c ~/.stackdio/supervisord.conf
supervisorctl -c ~/.stackdio/supervisord.conf start all
```

2.10 10. Try it out!

At this point, you should have everything configured and running, so fire up a web browser and point it to your hostname and you should see the stackd.io login page. If you're using LDAP, try logging in with a user that is a member of the `stackdio-admin` and `stackdio-user` groups, or login with the admin user you created earlier.

2.11 11. Creating additional users

Note: If you're using LDAP, you can skip this step.

The superuser we created earlier will give us admin access to stackd.io, however, you probably want at least one non-superuser. Point your browser to http://hostname:8000/__private/admin and use the username and password for the super user you created earlier. You should be presented with the Django admin interface. To create additional users, follow the steps below.

- click Users
- click Add user in the top right of the page
- set the username and password of the user and click save
- optionally provide first name, last name, and email address of the user and click save

The newly created users will now have access to stackd.io. Test this by logging out and signing in with one of the non-admin users.

Using the Amazon AMI

To make installation easier, we provide an `packer` build that creates an Amazon Machine Image (AMI). This AMI is built from an Ubuntu 14.04 LTS image, and is HVM based. While our ultimate plan is to provide this AMI on the [AWS Marketplace](#), we don't have this set up yet.

3.1 Building an AMI

Note: The build script that runs packer requires you to have python installed locally.

3.1.1 1. Install Packer

If you haven't already, install packer using their documentation [here](#). We recommend using homebrew for the installation if you're using OSX.

3.1.2 2. Accept License

Before building with packer, you must accept the license agreement for the base Ubuntu AMI: <http://aws.amazon.com/marketplace/pp?sku=b3dl4415quatdnd14qa6kcu45>

3.1.3 3. Clone Repository

```
git clone https://github.com/stackdio/stackdio.git
cd stackdio
```

3.1.4 4. Export AWS Credentials

Ensure packer knows about your aws credentials:

```
export AWS_ACCESS_KEY='<YOUR_ACCESS_KEY>'
export AWS_SECRET_KEY='<YOUR_SECRET_KEY>'
```

3.1.5 5. Run the packer build

Finally, run the packer build, where `<version>` is the version you want to build:

```
./packer/build.py <version>
```

After a few minutes, you should have a usable AMI.

3.2 Using the AMI

After you've built the AMI, you can launch an instance from it. Once the instance is running, you can navigate to `http://<instance-ip>/` and login using the following credentials:

```
username: admin  
password: stackdio
```

We recommend changing this password immediately after logging in the first time.

LDAP Guide

4.1 django-auth-ldap

Under the hood, we use `django-auth-ldap` for all our interaction with an LDAP server. It's a very useful library that allows us to integrate with LDAP in just a few lines of python configuration.

4.2 LDAP and stackd.io

If you'd like to integrate with LDAP, all you need to do is create an `ldap_settings.py` file in the `stackdio.server.settings` package. If this file exists, `stackd.io` will automatically pick it up and begin authenticating users to your ldap server. Inside that settings package, there is an `ldap_settings.py.template` file that you may rename to `ldap_settings.py` and update to match your needs. It's contents are displayed below.

The template contains most of the relevant bits you may want to use, but we defer anything deeper to the `django-auth-ldap` documentation.

Note: If you're using the packer-built AMI described in [Using the Amazon AMI](#), your ldap settings should be placed here: `/usr/share/stackdio/lib/python2.7/site-packages/stackdio/server/settings/ldap_settings.py`.

```
##
# LDAP configuration
# We are using django-auth-ldap to enable stackd.io to use LDAP for
# authentication. The settings below correspond to our internal
# LDAP and we don't guarantee this to work for all. Please see
# the docs at http://pythonhosted.org/django-auth-ldap/ for more
# information.
##

import ldap
from django_auth_ldap.config import LDAPSearch, GroupOfNamesType

# We use direct binding with a dedicated user. If you have anonymous
# access available or can bind with any user, you'll want to change
# this.
AUTH_LDAP_BIND_AS_AUTHENTICATING_USER = False
AUTH_LDAP_SERVER_URI = 'YOUR_LDAP_SERVER_URI'
AUTH_LDAP_BIND_DN = 'uid=YOUR_BIND_USER,ou=People,dc=yourcompany,dc=com'
AUTH_LDAP_BIND_PASSWORD = 'YOUR_BIND_USER_PASSWORD'
```

```
AUTH_LDAP_REQUIRE_GROUP = ('cn=stackdio-user,ou=Group,dc=yourcompany,dc=com')
AUTH_LDAP_USER_SEARCH = LDAPSearch('ou=People,dc=yourcompany,dc=com',
                                   ldap.SCOPE_SUBTREE,
                                   '(&(objectClass=Person)(uid=%(user)s))')

# Group handling. Read the django_auth_ldap documentation for more info.
AUTH_LDAP_FIND_GROUP_PERMS = False
AUTH_LDAP_MIRROR_GROUPS = False

AUTH_LDAP_GROUP_TYPE = GroupOfNamesType()
AUTH_LDAP_GROUP_SEARCH = LDAPSearch(
    'ou=Group,dc=yourcompany,dc=com',
    ldap.SCOPE_SUBTREE,
    '(objectClass=groupOfNames)'
)

AUTH_LDAP_USER_ATTR_MAP = {
    'first_name': 'givenName',
    'last_name': 'sn',
    'email': 'mail',
}

AUTH_LDAP_USER_FLAGS_BY_GROUP = {
    'is_superuser': 'cn=stackdio-admin,ou=Group,dc=yourcompany,dc=com',
    'is_staff': 'cn=stackdio-admin,ou=Group,dc=yourcompany,dc=com',
}

AUTH_LDAP_CONNECTION_OPTIONS = {
    ldap.OPT_X_TLS_REQUIRE_CERT: ldap.OPT_X_TLS_NEVER,
    ldap.OPT_X_TLS_NEWCTX: 0,
}
```

Webserver Guide

This guide will help you quickly get the web portion of stackd.io running behind either Apache or Nginx. You should've already worked through this manual install guide before running through the steps below. As with this guide, our focus is not entirely on building out a production-ready system, but merely helping you quickly get a system stood up to become familiar with stackd.io. Once you understand how it works, then we can start hardening the system for production use.

So, with that said, there are two paths to take: Apache or Nginx. We recommend using whichever you feel more comfortable with. Some of us here like Apache, while others like Nginx. Your mileage may vary :)

5.1 Common Steps

To do some of the steps below you will need to have already installed stackdio and be in the virtual environment. To make sure you're in the virtualenv:

```
workon stackdio
```

Both Apache and Nginx installs need a place to store logs and some static files to serve up. These two steps are common to both and should be run before proceeding with configuring the web server.

```
# Create the log directory that we'll have Apache or Nginx log to
mkdir -p ~/.stackdio/var/log/web

# And tell Django to collect its static files into a common directory for the webserver to serve up
stackdio manage.py collectstatic --noinput
```

5.2 Apache

5.2.1 CentOS Installation

Install required packages:

```
sudo yum install httpd mod_wsgi mod_ssl
```

Followed by having stackd.io generate a simple Apache configuration file for serving up the Django-based API and static assets and store the output into the appropriate location.

```
stackdio config apache | sudo tee /etc/httpd/conf.d/stackdio.conf > /dev/null
```

Fix a permissions problem with the user's home directory not having execute permissions. This is needed because of httpd v2.2 needing directory execute permissions from the web directory up to the root directory.

```
chmod +x ~/
```

Note: You may pass `--with-ssl` to generate boilerplate for serving over SSL, but you will need to add your certs and point to them in the configuration file. You may also need to remove the existing `ssl.conf` from within `conf.d`.

And that's it...let's start the server and then point your browser to the hostname on port 80 (use https if you decided to serve over SSL.)

```
sudo service httpd restart
```

5.2.2 Ubuntu Installation

Install required packages:

```
sudo apt-get install apache2 libapache2-mod-wsgi
```

and just like the CentOS instructions, generate and store the Apache configuration file into the correct location:

```
stackdio config apache | sudo tee /etc/apache2/sites-enabled/stackdio.conf > /dev/null
```

Note: You may pass `--with-ssl` to generate boilerplate for serving over SSL, but you will need to add your certs and point to them in the configuration file.

and finally, start the server:

```
sudo service apache2 restart
```

5.3 Nginx

In our configuration, Nginx will be used to serve static files and as a proxy to send requests down to the Django application running via gunicorn on port 8000. The configuration we'll generate is useful to use a quick start mechanism to get you up and running behind Nginx/gunicorn very quickly.

5.3.1 CentOS Installation

Install required packaged, generate and write configuration file, and restart server:

```
sudo yum install nginx

stackdio config nginx | sudo tee /etc/nginx/conf.d/stackdio.conf > /dev/null

# rename the default server configuration
sudo mv /etc/nginx/conf.d/default.conf /etc/nginx/conf.d/default.conf.bak

sudo service nginx restart
```

5.3.2 Ubuntu Installation

```
sudo apt-get install nginx

stackdio config nginx | sudo tee /etc/nginx/sites-enabled/stackdio.conf > /dev/null

# remove the default configuration symlink
sudo rm /etc/nginx/sites-enabled/default

sudo service nginx restart
```

Contact Information

- Static site - <http://stackd.io>
- IRC - [#stackdio](#) on Freenode
- Twitter - [@stackdio](#) or [#stackdio](#)
- Github - <https://github.com/stackdio/stackdio>

6.1 Filing issues and feature requests

We want stackd.io to be solid, but there's always going to be issues to fix and new features to enhance the product. If you find any problems or would like to request a particular feature, feel free to head over to [stackd.io's issue tracker](#). Below are some general guidelines.

6.2 Feature Requests

We definitely want this project to be better in every way, but sometimes we just can't see the forest because of those trees. If you think of something that would make stackd.io be more intuitive, easier, faster, or efficient, we encourage you to [file a feature request](#), but please keep the following in mind:

1. Double-check that the feature hasn't already made its way into the tracker and/or been completed
2. After filing it, please be patient – sometimes it might take a while for things to get prioritized, so don't be bummed that it's not available tomorrow ;)
3. As always, the quickest way to get your shiny new feature is to get your hands dirty and implement it yourself!

6.3 Bug Reports

If you think you've found a bug, here's what we ask of you:

1. Check the [issue tracker](#) and/or [search for your issue](#) to make sure we aren't already tracking the problem
2. See if you can reproduce the issue on the develop branch of the project to confirm we haven't already fixed the problem
3. Try to be as descriptive as possible when filing your issues. We'd like to know the operating system you're on, any stack traces you've seen in the logs, a good account of the steps required to reproduce the problem along

with what you feel is the expected outcome. Providing more information is almost always better so we can quickly identify and fix the issue.

4. Be patient :)
5. Optionally, it would be fantastic if you could help us out and fix the problem yourself...if you're brave enough, see the [Contributor Guidelines](#) on contributing to the project.

6.4 Usage Questions

Usage questions are best handled in one of the following ways:

- Asking a question on the IRC channel
- Emailing the contributors, especially those who have been recently active on the project.

6.5 Long term development comments

Long term development tasks will usually have a branch opened on the main repository, along with a pull request back into develop. This is the place to make comments on the development cycle for that feature.

Contributor Guidelines

We're hopeful that you'll find value in using stackd.io along with your existing tools to manage your infrastructure. We're also hopeful that you'll find ways to help contribute. Either through finding and reporting bugs, providing new features, or even getting your hands dirty and contributing some code or docs. However you feel comfortable contributing, we offer a few helpful guidelines to make it that much easier.

Note that stackd.io is built on SaltStack, and we believe that its community will find stackd.io useful. As such, we're trying to stay close to the conventions and guidelines they've adopted to make it easier for folks in that community to help out – and we've borrowed from the [SaltStack Development guide](#) :)

7.1 Filing issues and feature requests

The process for filing issues and feature requests is described in the [Contact Information](#) page.

7.2 Contributing Code

Since we're using Github, the recommended workflow for fixing bugs, adding features, or modifying documentation is to fork and submit pull requests. The process is pretty straightforward, but if you're unfamiliar with Github, take some time to browse through [Github's Help](#).

In a nutshell, we'll need you to:

- fork the stackd.io project into your personal account [[Tutorial](#)]
- make the necessary changes to the code/docs and issue a pull request. [[Tutorial](#)]
- keep your local fork in sync with the parent stackd.io repository to minimize the chance of merge conflicts. [[Tutorial](#)]
- and, if you're working on multiple things or your changes are going to be somewhat large, it's generally recommended to create a branch for each piece of work you're doing. [[Tutorial](#)]

NOTE: SaltStack has a [great guide](#) on how to work within their project and it mostly applies to stackd.io as well

7.3 Pull request guidelines

7.4 CLA

Contribution to stackd.io requires a CLA before pull requests will be merged. This is currently handled manually by the repo admins, but may be handled by a bot in the future.

7.5 Branch naming

The branch name that the pull request originates from should start with either `feature/` or `bugfix/`, depending on its contents. The rest of the branch name should describe the contents of the patch, preferably by being an Issue#. Issue#'s are required for `bugfix/` branches.

7.6 Code style and quality

7.7 PEP8 compatibility

All pull requests must meet PEP8 compatibility

7.8 Tests

Pull requests will be easier to review and understand if they contain automated tests for the functionality changed. As such, pull requests with tests are more likely to be accepted more quickly.