Kong

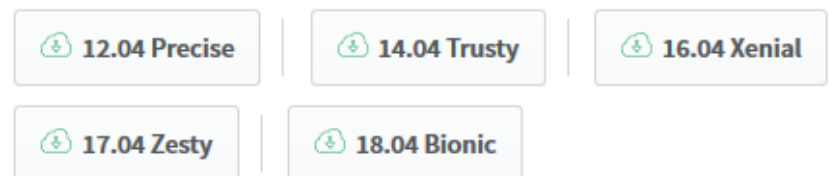# Ubuntu Installation

Table of Contents

- Packages
- APT Repositories
- Installation

Packages

Start by downloading the corresponding package for your configuration:

⬇ 12.04 Precise	⬇ 14.04 Trusty	⬇ 16.04 Xenial
⬇ 17.04 Zesty	⬇ 18.04 Bionic	

Enterprise trial users should download their package from their welcome email and save their license to `/etc/kong/license.json` after step 1.

APT Repositories

PACKAGES

- Docker
- Amazon Linux
- CentOS
- Red Hat
- Debian
- Ubuntu
- macOS

CLOUD

- DC/OS
- Kubernetes
- Google Cloud
- AWS Marketplace
- AWS Cloud Formation

SOURCE

- Vagrant
- Compile

You can also install Kong via APT; follow the instructions on the "Set Me Up" section on the page below, setting *distribution* to the appropriate value (lsb_release -sc) (e.g., `precise`) and *components* to `main` .

⬦ **Deb Repository**

Installation

① **Install Kong**

If you are downloading the [package](#), execute:

```
$ sudo apt-get update
$ sudo apt-get install openssl libpcre3 procps perl
$ sudo dpkg -i kong-1.3.0.*.deb
```

If you are using the apt repository execute:

```
$ sudo apt-get update
$ sudo apt-get install -y apt-transport-https curl lsb-core
$ echo "deb https://kong.bintray.com/kong-deb `lsb_release -sc
$ curl -o bintray.key https://bintray.com/user/downloadSubjectl
$ sudo apt-key add bintray.key
$ sudo apt-get update
$ sudo apt-get install -y kong
```

② **Prepare your database or declarative configuration file**

Kong can run either with or without a database.

When using a database, you will use the `kong.conf` configuration file for setting Kong's configuration properties at start-up and the database as storage of all configured entities, such as the Routes and Services to which Kong proxies.

When not using a database, you will use `kong.conf` its configuration properties and a `kong.yml` file for specifying the entities as a declarative configuration.

Using a database

Configure Kong so it can connect to your database. Kong supports both PostgreSQL 9.5+ and Cassandra 3.x.x as its datastore.

If you are using Postgres, please provision a database and a user before starting Kong, ie:

```
CREATE USER kong; CREATE DATABASE kong OWNER kong;
```

Now, run the Kong migrations:

```
$ kong migrations bootstrap [-c /path/to/kong.conf]
```

Note for Kong < 0.15: with Kong versions below 0.15 (up to 0.14), use the `up` sub-command instead of `bootstrap`. Also note that with Kong < 0.15, migrations should never be run concurrently; only one Kong node should be performing migrations at a time. This limitation is lifted for Kong 0.15, 1.0, and above.

Without a database

If you are going to run Kong in DB-less mode, you should start by generating declarative config file. The following command will generate a `kong.yml` file.

declarative config file. The following command will generate a `kong.yml` file in your current folder. It contains instructions about how to fill it up.

```
$ kong config init
```

After filling up the `kong.yml` file, edit your `kong.conf` file. Set the `database` option to `off` and the `declarative_config` option to the path of your `kong.yml` file:

```
database = off
declarative_config = /path/to/kong.yml
```

3   Start Kong

```
$ kong start [-c /path/to/kong.conf]
```

4   Use Kong

Kong is running:

```
$ curl -i http://localhost:8001/
```

Quickly learn how to use Kong with the 5-minute Quickstart.

✉   Follow Up:

Get future release notes emailed to you.

Email

Notify me

Having Trouble? We're Here to Help!

Kong

[Email] [Subscribe]

PRODUCTS

Kong

Kong Enterprise

Kong Cloud

Subscriptions

RESOURCES

Enterprise Support

Documentation

Partners

Webinars

Ebooks

COMPANY

About

Investors

News

Careers HIRING!

Kong Summit

Contact

OPEN SOURCE

Install

GitHub

Kong Nation

Community

© Kong Inc. 2019   **Terms**   ·   **Privacy**