
cinch Documentation

Release

RedHatQE

Apr 19, 2017

Contents

1	Users	1
2	Development	7
3	Docker Image	9
4	User Files	11
5	Indices and tables	13

Quick Start

If you would rather not be bothered with the boring, gritty details, the following quick steps will get you up and going with a Cinch installation, after which you can jump down to the section on how to run Cinch, and skip the installation paragraphs that immediately follow this one.

Fedora 25+

Execute the following commands

```
sudo dnf install -y libvirt-devel python-virtualenv libyaml-devel  
openssl-devel libffi-devel gcc redhat-rpm-config  
virtualenv cinch && source cinch/bin/activate  
pip install cinch
```

RHEL/CentOS

See the sample Ansible playbook for steps to install Cinch into a virtualenv in RHEL/CentOS, as there are some additional steps required in these systems. Running that playbook as-is on your local system will result in a virtualenv living at `/var/lib/jenkins/opt/cinch` that will contain the latest version of Cinch.

Ubuntu

```
apt-get install -y libvirt-dev python-virtualenv libyaml-dev openssl  
libffi-dev gcc python-dev libssl-dev  
virtualenv cinch && source cinch/bin/activate
```

On older versions of Ubuntu (like 14.04) you should update pip (usually these older systems come with a version of pip such as 1.5.4 instead of version 9+)

```
pip install -U pip
```

On all systems, continue on with the installation of Cinch itself

```
pip install cinch
```

Note: After Cinch is installed with the above quick start methods, you can jump down to the section about running Cinch to get a look at basic documentation on how the software should be used.

Getting Started

At its core, this software is nothing more than a collection of Ansible playbook scripts for configuring a system. Any knowledge that you have that is applicable to the broad spectrum of Ansible usage is applicable here. You can opt to install Ansible from your favorite package manager, or you can use the version that is specified in **setup.py**

Before concluding there is a bug in these playbooks, make sure that the version of Ansible you are using is the same as the version in the **setup.py** file and that you have ensured there are no alterations from that version. It is not intended or guaranteed that any changes from the stock version of Ansible that has been tested should work.

Requirements

To setup your environment, you need to install Ansible. Since Ansible is primarily distributed as a Python package, it is suggested that you use **pip** to install Ansible on your system. You are welcome to try and use the version that is installed by your favorite package manager, but be sure that you are using a version at least as new as the version pinned in **setup.py**.

It is recommended that you install Ansible from **pip** using a virtualenv, as is the best practices recommendations for most Python packages that are available from PyPI. In order to build and install Ansible, you will need to install the following system packages:

Note: If you install cinch via **pip**, a supported version of Ansible will be brought in as a dependency.

- gcc or appropriate system compiler
- OpenSSL development package
- libyaml development package
- virtualenv package
- libffi development package
- libvirt development package

Use your system package manager to install these packages, if they are not already present.

Note: You will need to install the development version of the libraries, as **pip** will attempt to build wrappers around some of those libraries during its install of Ansible and dependencies.

Here is an example of installing required system level packages for Fedora 25:

Installation

Once the system level packages are installed, you can install cinch using **pip** (virtualenv strongly recommended):

Fedora

```
virtualenv cinch && source cinch/bin/activate
pip install cinch
```

RHEL7 and CentOS7

RHEL7 and derivatives offer older versions of Python packaging tools that are incompatible with some cinch dependencies. To work-around this issue, we have provided an [Ansible playbook](#) that will install a newer version of the necessary Python packaging tools to allow for installation on RHEL7. This playbook is intended for use on Jenkins masters and will install cinch into a virtualenv at **/var/lib/jenkins/opt/cinch**. For convenience, an optional [Jenkins Job Builder template](#) is provided and will create a Jenkins job that will run the aforementioned playbook on your Jenkins master.

Execution

With linch-pin

The `cinchpin` command can be used to call [linch-pin](#) automatically to provision instances and then configure the instances. `cinchpin` supports a subset of linch-pin commands, such as **rise**, **drop**, and **init**.

In the following example we will provision a RHEL7 instance in OpenStack as a Jenkins slave.

First, create necessary credentials for linch-pin provisioning for your target infrastructure in `<venv-path>/lib/python2.7/site-packages/provision/roles/openstack/vars/os_creds.yml`:

```
---
# openstack API endpoint
endpoint: http://openstack-api-endpoint.example.com:5000/v2.0

# project/tenant name
project: myproject

# project/tenant username and password
username: myuser
password: mypass
```

Note: The upcoming auth driver feature for linch-pin will make this step easier in the future.

Next, generate a linch-pin working directory for use with cinch by running the following commands:

```
mkdir /path/to/workdir
cinchpin init -w /path/to/workdir
```

Create a layout file by saving the following example template as `/path/to/workdir/layouts/cinch.yml` and edit to taste. For the `jenkins_user_password` variable, please use the [Ansible documentation](#) to generate a suitable password hash. **For security in production environments, DO NOT copy the existing hash from this example.**

```
---
inventory_layout:
  hosts:
    cinch:
      count: 1
      # List all necessary 'cinch' roles here
      host_groups:
        - rhel7 # specify the appropriate Ansible role for your distribution
        - certificate_authority # optional role to install CA certificates
        - repositories # the 'repositories' role is required for a Jenkins master or
↪ slave
        - jenkins_slave
      host_groups:
        all:
          vars:
            # required variables for all hosts
            ansible_user: root
            ansible_private_key_file: /path/to/ssh/private_key
            ansible_connection: ssh
          certificate_authority:
            vars:
              # Add URLs from which to download CA certificates for installation
            certificate_authority_urls:
              - https://example.com/ca1.crt
              - https://example.com/ca2.crt
          repositories:
            vars:
              # Base URL for repository mirror
            rhel_base: http://example.com/content/dist/rhel/server/7/7Server
          jenkins_slave:
            vars:
              # Required variables for a Jenkins slave

              # The password for the Jenkins user account that will be created on the slave.
              # For security in production environments, DO NOT copy the
              # existing hash from this example.
              # https://docs.ansible.com/ansible/faq.html#how-do-i-generate-crypted-
↪ passwords-for-the-user-module
            jenkins_user_password: '$6$rounds=656000$YQKMBktZ/Gaggxf0
↪ $KC7xhatWzdDJyvCDo7htomtISsvd2MWN87RB3TsAbq1Nmwwdy/z2Et8kQil/
↪ tZkHjfd2vGlr7W2R9rjpaA1C5/'
            jenkins_master_url: 'http://jenkins.example.com' # URL to Jenkins master for
↪ the slave to connect to
```

Create a topology file by saving the following example template as `/path/to/workdir/topologies/cinch.yml` and edit to taste:

```
---
topology_name: "cinch_topology"

# OpenStack project/tenant name
site: "my-openstack-project-name"

resource_groups:
```



```

-
  resource_group_name: "cinch"
  res_group_type: "openstack"
  res_defs:
    - res_name: "resource"
      flavor: "m1.large"
      res_type: "os_server"
      image: "rhel-7.2-server-x86_64-released"
      count: 1 # Number of instances to create
      keypair: "openstack-keypair-name" # Name of SSH keypair configured for
↪OpenStack account
      networks:
        - "openstack-network-name" # OpenStack network name

      # Name of credentials file to use for the OpenStack API
      assoc_creds: "os_creds"

```

Note: For more topology examples, including various host environments, see the [linch-pin documentation](#).

Provision and configure your Jenkins slave automatically with the following command:

```
cinchpin rise -w /path/to/workdir
```

To terminate the OpenStack instance and remove the Jenkins slave from the Jenkins master, run the following command:

```
cinchpin drop -w /path/to/workdir
```

Note: Once the working directory is configured successfully, a common next step would be to check this directory into source control where it can be consumed by CI automation tools such as Jenkins Job Builder or Jenkins Pipeline.

Manual

Execution of this software requires configuring an Ansible inventory that points at the **jenkins_master** and **jenkins_slave** hosts that you want configured. Use normal methods for setting **group_vars** and **host_vars** within the inventory or its associated folders that suits your own needs and preferences.

While most default settings should be functional, there are lots of options configured in the various **default/main.yml** files within the various roles folders. Check in those files for more details on specific options that can be set and a description of what they each mean.

See a few examples of such in either the **inventory/** folder or inside of the various **vagrant/** subfolders where known good working environments are configured for development use.

The path **inventory/local** is excluded from use by the project and can be leveraged for executing and storing your own local inventories, if the desire arises. There is even a shell script in **bin/run_jenkins_local.sh** that will execute **ansible-playbook** from the **.venv/** virtualenv and point it to the **inventory/local/hosts** file to make executing against your own environment as easy as a single command.

Support

The playbooks should support, minimally, CentOS and RHEL versions 7+. If you encounter difficulties in those environments, please file bugs. There should be no configuration necessary for a CentOS host, and a RHEL host requires only that you configure the base URL for your local RHEL repository collection. See documentation in the appropriate roles for details on that configuration.

Environments

Development occurs targeting each of the specific host environments that are supported. The default development environment and targeted host is the latest version of CentOS.

The fastest way to get yourself up and running is to leverage the Vagrant machines held within the top-level vagrant folder. These are named according to the roles that each one is designed to exercise.

Install

To run the software locally, you need a few basic pieces of software installed. The following packages for Fedora need to be installed, minimally, or the equivalent packages for your distribution:

- python-virtualenv
- gcc
- redhat-rpm-config
- openssl-devel
- libvirt-devel
- libyaml-devel
- vagrant

The only software actually required to run the playbooks is Ansible and its dependencies. The other packages listed above are required only to install and build Ansible and its dependencies, such as PyYAML. Thus, if you are looking to package Cinch for a new distribution, the above packages, less vagrant, are a good starting place for build dependencies.

If installing manually, you can activate your Python virtualenv of choice and issue the command `pip install /path/to/cinch`. As a developer, if you plan to make changes to Cinch, then use pip in the local editable mode by issuing the command `pip install -e /path/to/cinch` instead.

Execution

Once all of these dependencies are fulfilled, there are a number of folders under the top level `vagrant/` directory that contain minimally a Vagrantfile. The Vagrantfile can be used to issue the command “`vagrant up`” from within that directory to spin up a collection of machines, against which the cinch playbooks will be automatically executed. Consult the README in each directory for more information about which machines will be created out of that directory, and for any information that the user might need to supply.

Some of the Vagrantfile values will need to be supplied by the user, specifically any values related to RHEL repository URLs as there is no public version of those repositories available. Other values should all be provided from within those directories already.

Merely issuing the command `vagrant up` should bring up the VMs for each environment you configure. For the most part, it should be possible to run each environment on your local system, but there is the potential that having multiple environments running at the same time on the same host could result in collisions between the IP addresses of the hosts. It certainly would lead to provided URLs in the README files being incorrect.

Docker Image

For users who do not want to provision an entire system to run a Jenkins slave there exists a Docker image which can quickly get a Jenkins Swarm connected instance to run.

Source Image

For every release of cinch that is made, a version of the Docker container is pushed to Docker Hub. Multiple tags are pushed for each Cinch release. They are named by combining source information image along with the version of Cinch used to build them.

Currently there are images built off of

- centos:7
- centos:6

These get tagged into the Cinch image repository as

- redhatqecinch/jenkins_slave:cent7-0.5.2
- redhatqecinch/jenkins_slave:cent6-0.5.2

This indicates two images, one based on the centos:7 image and one based off the centos:6 image. Both of them are built by the version 0.5.2 release of Cinch.

Image Options

As with the rest of Cinch, there are some customizable image options that a user must supply before the image will work with your infrastrucutre. However, unlike using the Ansible-based solution to create your own system, there are far fewer options. Other than the following options, all builds of the Cinch Docker images utilize all default values for a Cinch slave instance.

There are two variables that the user is required to provide before the image will run properly. Those are

Environment Variable	Explanation
JENKINS_MASTER_URL	The URL to the Jenkins master instance that this slave should connect to
JSLAVE_NAME	The name this slave will be given on the Master node
JSLAVE_LABEL	The Jenkins label this slave will receive, which will be matched against jobs requiring certain labels for execution
JSWARM_EXTRA_ARGS	Additional command-line arguments to pass to the JSwarm client in the image

If the container image is run directly from the Docker command line, these options may be passed through *docker*'s `-e` option. When running the image in Kubernetes or OpenShift, use that system's methods for passing in environment variables to the image.

Customizing the Image

Instead of running the base image provided, a group could choose to use a Dockerfile to extend the base image provided to do such things as install custom software, edit configurations, etc. If that is the case, then the environment variables can absolutely be preset within the Dockerfile using its `ENV` command, as with any other environment variable.

Extending the image in this way could simplify deployment, as the image could include information such as the Jenkins Master URL already configured to connect to the organization's Jenkins instance. Likewise, different slave images could be pre-populated with packages and slave labels for building different types of software or running different types of tasks. As nothing more than a standard Docker image, the provided images can be made fully extensible.

One note is that the image is set to run all commands as the user "jenkins". If the image is being extended, then it might be necessary to set the `USER` command in the extending Dockerfile to "USER root" if system software is being installed.

Motivation

Anyone using Cinch to provision either a Jenkins master or slave may have the need to perform configuration to the system that exceeds the ability of Cinch to reasonably include support for within these playbooks. These could cover nearly any aspect of system administration, monitoring, configuration, and setup. For such a case, it is recommended that the user leverage the ability of Ansible to file a host into multiple different inventory groups, and private configuration be stored in private playbooks. Then those playbooks can be executed either before or after (or both) the Cinch playbooks are executed.

However, there are a few basic system administration tasks that are general enough, and simple enough, that Cinch has opted to support those features to assist in the configuration of a Jenkins master. In addition to supporting the ability to setup Yum/DNF repositories during configuration and configure certificate authority chains, both of which are important to installing the packages required by Cinch and to configure SSL options for Jenkins, another feature supported by Cinch is the ability to upload arbitrary files from the local system where Ansible is being hosted to the remote system being configured.

Mechanisms

Each Ansible host, or group, can have defined values of files to upload to the remote hosts. These uploads happen at two different points during the execution of Cinch. The first set of uploads occurs before any Cinch plays have been executed except for verifying the host is reachable. This means that none of the Cinch-related configurations will be available during this upload run, unless they have previously been configured. This includes things like the “jenkins” system user, configured repositories, certificate authorities, etc. The second run happens at the very end - after both the master and any slaves have been configured and are up and running. However, at this point, all such configurations, users, etc are already present on the system.

Thus, it is important to realize a file cannot be uploaded to be owned by the Jenkins user before the Jenkins user is created. If it is necessary to upload a file as that user before the Jenkins service starts on a configured host, then it will be necessary to use external playbooks or other methods to ensure proper behavior.

Configuration

Configuring uploads either before or after a Cinch run is straightforward. Simply override the values of the arrays “pre_upload_files” and “post_upload_files” in the Ansible host or group configurations for all hosts that require such a feature.

These arrays require identical structures. Each element in the array should be an object hash with certain values defined. Those values are listed below:

value	required?
src	yes
dest	yes
owner	no
group	no
mode	no

Example:

```
pre_upload_files:
- src: /home/deployuser/somehost/ssl.key
  dest: /etc/apache2/ssl/ssl.key
  mode: 0600
post_upload_files:
- src: /home/deployuser/somehost/ssh
  dest: /var/lib/jenkins/.ssh
  owner: jenkins
  mode: 0600
```

Each of these values is passed directly into the Ansible module called `copy`. Refer to that module’s documentation for information about the structure and values that are permitted to be passed into these values. Note, especially, that this module can be used to upload whole directories in addition to individual files.

If the need arises to support more of the options of that module, adding that support to Cinch can be done. Please just open an issue in the [GitHub Issue Tracker](#) detailing the requested functionality.

CHAPTER 5

Indices and tables

- `genindex`
- `search`