
LinchPin Documentation

Release 1.0.0

Samvaran Kashyap Rallabandi

Jun 01, 2017

Contents

1	About LinchPin	1
1.1	Introduction	1
1.2	Installation	1
1.3	Getting Started	3
1.4	Configuration	8
1.5	Python API Reference	15
1.6	Glossary	19
2	Indices and tables	23
	Python Module Index	25

CHAPTER 1

About LinchPin

Welcome to the LinchPin documentation!

LinchPin is a hybrid cloud orchestration tool. Its intended purpose is managing cloud resources across multiple infrastructures. These resources can be provisioned, decommissioned, and configured all using a topology file and a simple command-line interface.

Additionally, LinchPin provides a Python API (and soon a RESTful API) for managing resources. The cloud management component is backed by *Ansible* <<https://ansible.com>>. The front-end API manages the interface between the command line (or other interfaces) and calls to the Ansible API.

This documentation covers the current released version of LinchPin (1.0.0). For recent features, we attempt to note in each section the version of LinchPin where the feature was added.

Introduction

Before getting heavily into LinchPin, let's cover some of the basics. The topics below will cover everything needed to get going with LinchPin. For more advanced topics, refer to the [main documentation page](#).

- [Installation](#)
- [Getting Started](#)
- [Configuration](#)

See also:

User Mailing List Subscribe and participate. A great place for Q&A

irc.freenode.net #linchpin IRC chat channel

Installation

Topics

- *Installation*
 - *Minimal Software Requirements*
 - *Installing LinchPin*
 - * *Source Installation*

Currently, LinchPin can be run from any machine with Python 2.6+ (Python 3.x is currently experimental), and requires Ansible 2.2.1. There are many other dependencies, depending on the provider. The core providers are *OpenStack*, *Amazon EC2*, and *Google Compute Cloud*. If enabled on the host system, *Libvirt* can also be used out of the box.

Refer to your specific operating system for directions on the best method to install Python, if it is not already installed. Many modern operating systems will have Python already installed. This is typically the case in all versions of Linux and OS X, but the version present might be older than the version needed for use with Ansible. You can check the version by typing `python --version`.

If the system installed version of Python is older than 2.6, many systems will provide a method to install updated versions of Python in parallel to the system version (eg. `virtualenv`).

Minimal Software Requirements

As LinchPin is heavily dependent on Ansible, this is a core requirement. Beyond installing Ansible, there are several packages that need to be installed:

```
* libffi-devel
* openssl-devel
* libyaml-devel
* gmp-devel
* libselinux-python
* make
```

For Fedora/CentOS/RHEL the necessary packages should be installed.

```
$ sudo yum install python-virtualenv libffi-devel \
openssl-devel libyaml-devel gmp-devel libselinux-python make
```

Note: Fedora will present an output suggesting the use of *dnf* as a replacement for yum.

Installing LinchPin

Note: Currently, linchpin is not packaged for any major Operating System. If you'd like to contribute your time to create a package, please contact the [linchpin mailing list](#).

Create a virtualenv to install the package using the following sequence of commands (requires `virtualenvwrapper`).

```
$ mkvirtualenv linchpin
..snip..
(linchpin) $ pip install linchpin
..snip..
```

To deactivate the virtualenv.

```
(linchpin) $ deactivate
$
```

Then reactivate the virtualenv.

```
$ workon linchpin
(linchpin) $
```

If testing or docs is desired, additional steps are required.

```
(linchpin) $ pip install linchpin[docs]
(linchpin) $ pip install linchpin[tests]
```

Source Installation

As an alternative, LinchPin can be installed via github. This may be done in order to fix a bug, or contribute to the project.

```
(linchpin) $ git clone git://github.com/CentOS-PaaS-SIG/linch-pin
..snip..
(linchpin) $ pip install ./linch-pin
```

See also:

[User Mailing List](#) Subscribe and participate. A great place for Q&A

[irc.freenode.net](#) #linchpin IRC chat channel

Getting Started

Topics

- *Getting Started*
 - *Foreword*
 - *Terminology*
 - * *Topology*
 - * *Inventory Layout*
 - * *PinFile*
 - *Running linchpin*
 - * *Initialization (init)*
 - * *Provisioning (up)*
 - * *Teardown (destroy)*
 - * *Multi-Target Actions*

Foreword

Now that LinchPin is installed according to [Installation](#), it is time to see how it works. This guide is essentially a quick start guide to getting up and running with LinchPin.

LinchPin is a command-line utility, a Python API, and Ansible playbooks. This document focuses on the command-line interface.

Terminology

LinchPin, while it attempts to be a simple tool for provisioning resources, still does have some complexity. To that end, this section attempts to define the minimal bits of terminology needed to understand how to use the `linchpin` command-line utility.

Topology

The *topology* is a set of rules, written in YAML, that define the way the provisioned systems should look after executing `linchpin`. Generally, the *topology* and *topology_file* values are interchangeable, except where the YAML is specifically indicated. A simple **dummy** topology is shown here.

```
---
topology_name: "dummy_cluster" # topology name
resource_groups:
-
  resource_group_name: "dummy"
  resource_group_type: "dummy"
  resource_definitions:
  -
    name: "web"
    type: "dummy_node"
    count: 3
```

This topology describes a set of three (3) dummy systems that will be provisioned when *linchpin up* is executed. The names of the systems will be `web_#.example.net`, where # indicates the count (usually 0, 1, and 2). Once provisioned, the resources will be output and stored for reference. The output *resources* data can then be used to generate an inventory, or passed as part of a *linchpin destroy* action.

Inventory Layout

The *inventory_layout* or *layout* mean the same thing, a YAML definition for providing an Ansible static inventory file, based upon the provided topology. A YAML *layout* is stored in a *layout_file*.

```
---
inventory_layout:
  vars:
    hostname: __IP__
  hosts:
    example-node:
      count: 3
      host_groups:
        - example
  host_groups:
    example:
```



```
vars:
  test: one
```

The above YAML allows for interpolation of the ip address, or hostname as a component of a generated inventory. A host group called *example* will be added to the Ansible static inventory, along with a section called *example:vars* containing *test = one*. The resulting static Ansible inventory is shown here.

```
[example:vars]
test = one

[example]
web-2.example.net hostname=web-2.example.net
web-1.example.net hostname=web-1.example.net
web-0.example.net hostname=web-0.example.net

[all]
web-2.example.net hostname=web-2.example.net
web-1.example.net hostname=web-1.example.net
web-0.example.net hostname=web-0.example.net
```

PinFile

A *PinFile* takes a *topology* and an optional *layout*, among other options, as a combined set of configurations as a resource for provisioning. An example *Pinfile* is shown.

```
dummy1:
  topology: dummy-cluster.yml
  layout: dummy-layout.yml
```

The *PinFile* collects the given *topology* and *layout* into one place. Many *targets* can be referenced in a single *PinFile*.

The *target* above is named *dummy1*. This *target* is the reference to the *topology* named *dummy-cluster.yml* and *layout* named *dummy-layout.yml*. The *PinFile* can also contain definitions of *hooks* that can be executed at certain pre-defined states.

Running linchpin

As stated above, this guide is about using the command-line utility, *linchpin*. First off, simply execute *linchpin* to see some options.

```
$ linchpin
Usage: linchpin [OPTIONS] COMMAND [ARGS]...

    linchpin: hybrid cloud orchestration

Options:
  -c, --config PATH          Path to config file
  -w, --workspace PATH       Use the specified workspace if the familiar Jenkins
                             $WORKSPACE environment variable is not set
  -v, --verbose              Enable verbose output
  --version                  Prints the version and exits
  --creds-path PATH         Use the specified credentials path if WORKSPACE
                             environment variable is not set
  -h, --help                 Show this message and exit.
```

```
Commands:
  init      Initializes a linchpin project.
  up        Provisions nodes from the given target(s) in...
  destroy   Destroys nodes from the given target(s) in...
```

What can be seen immediately is a simple description, along with options and arguments that can be passed to the command. The three commands found near the bottom of this help are where the focus will be for this document.

Initialization (init)

Running `linchpin init` will generate the directory structure needed, along with an example *PinFile*, *topology*, and *layout* files. One important option here, is the `-workspace`. When passing this option, the system will use this as the location for the structure. The default is the current directory.

```
$ export WORKSPACE=/tmp/workspace
$ linchpin init
PinFile and file structure created at /tmp/workspace
$ cd /tmp/workspace/
$ tree
.
- credentials
- hooks
- inventories
- layouts
|   - example-layout.yml
- PinFile
- resources
- topologies
  - example-topology.yml
```

At this point, one could execute `linchpin up` and provision a single libvirt virtual machine, with a network named *linchpin-centos71*. An inventory would be generated and placed in `inventories/libvirt.inventory`. This can be known by reading the `topologies/example-topology.yml` and gleaning out the *topology_name* value.

Provisioning (up)

Once a *PinFile*, *topology*, and optionally a *layout* are in place, provisioning can happen.

Note: For this section, the dummy tooling will be used as it is much simpler and doesn't require anything extra to be configured. The dummy provider creates a temporary file, which represents provisioned hosts. If the temporary file does not have any data, hosts have not been provisioned, or they have been recently destroyed.

The dummy *topology*, *layout*, and *PinFile* are shown above in the appropriate sections. The tree would be very simple.

```
$ tree
.
- inventories
- layouts
|   - dummy-layout.yml
- PinFile
- resources
- topologies
  - dummy-cluster.yml
```

Performing the command `linchpin up` should generate *resources* and *inventory* files based upon the *topology_name* value. In this case, is `dummy_cluster`.

```
$ linchpin up
target: dummy1, action: up

$ ls {resources,inventories}/dummy*
inventories/dummy_cluster.inventory  resources/dummy_cluster.output
```

To verify resources with the dummy cluster, check `/tmp/dummy.hosts`

```
$ cat /tmp/dummy.hosts
web-0.example.net
web-1.example.net
web-2.example.net
```

This is reflected in both the *resources* (not shown) and *inventory* files.

```
[example:vars]
test = one

[example]
web-2.example.net hostname=web-2.example.net
web-1.example.net hostname=web-1.example.net
web-0.example.net hostname=web-0.example.net

[all]
web-2.example.net hostname=web-2.example.net
web-1.example.net hostname=web-1.example.net
web-0.example.net hostname=web-0.example.net
```

Teardown (destroy)

As expected, LinchPin can also perform *teardown* of *resources*. A teardown action generally expects that resources have been *provisioned*. However, because Ansible is idempotent, `linchpin destroy` will only check to make sure the resources are up. Only if the resources are already up will the teardown happen.

The command `linchpin destroy` will either use *resources* and/or *topology* files to determine the proper *teardown* procedure. The *dummy* Ansible role does not use the resources, only the topology during teardown.

```
$ linchpin destroy
target: dummy1, action: destroy

$ cat /tmp/dummy.hosts
-- EMPTY FILE --
```

Note: The teardown functionality is slightly more limited around ephemeral resources, like networking, storage, etc. It is possible that a network resource could be used with multiple cloud instances. In this way, performing a `linchpin destroy` does not teardown certain resources. This is dependent on each providers implementation.

See specific implementations for each of the *providers*.

Multi-Target Actions

LinchPin can *provision* or *teardown* any number of *resources*. If a *PinFile* has multiple *targets*, and is called without a target name, all targets will be executed. Given this PinFile.

```
example:
  topology: example-topology.yml
  layout: example-layout.yml

example2:
  topology: example2-topology.yml
  layout: example2-layout.yml

dummy1:
  topology: dummy-cluster.yml
  layout: dummy-layout.yml
```

A call to `linchpin up` would *provision* and generate an Ansible static *inventory* for each *target*.

```
$ linchpin up
target: dummy1, action: up

target: example2, action: up

target: example, action: up
```

See also:

linchpincli Linchpin Command-Line Interface

User Mailing List Subscribe and participate. A great place for Q&A

irc.freenode.net #linchpin IRC chat channel

Configuration

Before resources can be provisioned in any of the environments through the use of linchpin, the environment must be configured to specify the resources required.

General Configuration

Managing LinchPin requires a few configuration files. Beyond `linchpin.conf`, there are a few other configurations that need to be created.

Topics

- *General Configuration*
 - *Workspace*
 - *Initialization*
 - *PinFile*
 - *Topologies*

Workspace

Initialization

Running `linchpin init` will generate the directory structure needed, along with an example *PinFile*, *topology*, and *layout* files. One important option here, is the `–workspace`. When passing this option, the system will use this as the location for the structure. The default is the current directory.

```
$ export WORKSPACE=/tmp/workspace
$ linchpin init
PinFile and file structure created at /tmp/workspace
$ cd /tmp/workspace/
$ tree
.
- credentials
- hooks
- inventories
- layouts
|   - example-layout.yml
- PinFile
- resources
- topologies
  - example-topology.yml
```

At this point, one could execute `linchpin up` and provision a single libvirt virtual machine, with a network named *linchpin-centos71*. An inventory would be generated and placed in `inventories/libvirt.inventory`. This can be known by reading the `topologies/example-topology.yml` and gleaning out the *topology_name* value.

PinFile

A *PinFile* takes a *topology* and an optional *layout*, among other options, as a combined set of configurations as a resource for provisioning. An example *Pinfile* is shown.

```
dummy1:
  topology: dummy-cluster.yml
  layout: dummy-layout.yml
```

The *PinFile* collects the given *topology* and *layout* into one place. Many *targets* can be referenced in a single *PinFile*.

The *target* above is named *dummy1*. This *target* is the reference to the *topology* named *dummy-cluster.yml* and *layout* named *dummy-layout.yml*. The *PinFile* can also contain definitions of *hooks* that can be executed at certain pre-defined states.

Topologies

The *topology* is a set of rules, written in YAML, that define the way the provisioned systems should look after executing `linchpin`. Generally, the *topology* and *topology_file* values are interchangeable, except where the YAML is specifically indicated. A simple **dummy** topology is shown here.

```
---
topology_name: "dummy_cluster" # topology name
resource_groups:
  -
    resource_group_name: "dummy"
    resource_group_type: "dummy"
    resource_definitions:
      -
        name: "web"
        type: "dummy_node"
        count: 3
```

This topology describes a set of three (3) dummy systems that will be provisioned when *linchpin up* is executed. The names of the systems will be ‘web_*#*.example.net’, where *#* indicates the count (usually 0, 1, and 2). Once provisioned, the resources will be output and stored for reference. The output *resources* data can then be used to generate an inventory, or passed as part of a *linchpin destroy* action.

Inventory Layouts

The *inventory_layout* or *layout* mean the same thing, a YAML definition for providing an Ansible static inventory file, based upon the provided topology. A YAML *layout* is stored in a *layout_file*.

```
---
inventory_layout:
  vars:
    hostname: __IP__
  hosts:
    example-node:
      count: 3
      host_groups:
        - example
  host_groups:
    example:
      vars:
        test: one
```

The above YAML allows for interpolation of the ip address, or hostname as a component of a generated inventory. A host group called *example* will be added to the Ansible static inventory, along with a section called *example:vars* containing *test = one*. The resulting static Ansible inventory is shown here.

```
[example:vars]
test = one

[example]
web-2.example.net hostname=web-2.example.net
web-1.example.net hostname=web-1.example.net
web-0.example.net hostname=web-0.example.net

[all]
web-2.example.net hostname=web-2.example.net
web-1.example.net hostname=web-1.example.net
web-0.example.net hostname=web-0.example.net
```

Topologies

A topology is a specification of which resources from which environments are being requested from a linchpin run. Since each environment has different sets of requirements, the exact values and structure of a topology file will vary based on where resources are to be provisioned. In this document some broad discussion of topologies will be addressed. More extensive examples pertaining to specific environments will be given in a separate section of the documentation.

Topology

Broadly speaking, a linchpin topology file is a list of resources to be provisioned from each environment. It is possible and a very common use case to list multiple resources, even multiple types of resources, in a single topology file. A less common use case, but still supported, is to provision multiple resources across multiple environments.

The topology file does not designate the format of the output, nor map the particular resources that get provisioned into output inventory groups. That is the work of the layouts file.

Structure

A topology is a YAML file or a JSON file formatted with a single top-level object.

There are two top level keys in a topology.

The first key is named *topology_name* and is a relatively free-form string that identifies the user-friendly name for this particular topology. For best practices, this should resemble the file name and possibly the name of the key from the PinFile.

The second key is the *resource_groups* key. This key is an array of objects.

Resource Group

Each entry in the *resource_group* key array is itself an object hash with three object keys.

The first key is *resource_group_name*, and is a user-friendly name for the resources that will be provisioned from this group definition.

The second key is *res_group_type* and must be a string of a limited set. This set must match to the particular environment. Some environments can define different types of valid values. As an example, the value *duffy* will define a resource type to be provisioned in a Duffy architecture, whereas the value *beaker* will contain definitions of a set of servers to be provisioned in a Beaker environment.

The third key is *res_defs*. This key defines an array of objects. Each of these objects' exact form will be dictated by the value of *res_group_type*. To see more information on the structure of these values, check the example topologies section of this documentation.

Layouts

A layout file is the current mechanism to define mappings between the resources provisioned out of the topology and the Ansible inventory groups that are output.

Topics

- [Layouts](#)

– *Structure*
* *Hosts*

Structure

As with a topology file, a layout file is a YAML file or a JSON file with a single root object hash. There is one top-level entry in the hash, named *inventory_layout*. The *inventory_layout* value is itself an object that has a few fields inside of it.

Hosts

The first hash value is *hosts*, which contains an object hash as a value. The keys of that hash are the names of hosts that have been provisioned out of the topology. Each host has two properties, *count* and *host_groups*.

The *count* property says how many of the topology hosts are to share this inventory hostname. For instance, if the host is “webserver” and *count* is 2, then this will generate hosts in the output inventory named “webserver-1” and “webserver-2”. This value is optional and defaults to 1 when it’s not present.

The *host_groups* field contains an array of Ansible inventory groups into which all the hosts under this hash will get placed. This value is optional and will default to an empty list when not filled. In that case, the host will be named into the inventory with its host vars, and added to default ‘all’ group.

As an example, assume you provisioned three hosts and you wanted one database and two applicaiton hosts. Your goal is to get to an Ansible inventory that looks like this:

```
[backend]
database

[frontend]
webhost-1
webhost-2

[ldap]
database
webhost-1
webhost-2

[security_updates]
database
```

Then your hosts object would look like this:

```
hosts:
  database:
    count: 1
    host_groups:
      - backend
      - ldap
      - security_updates
  webhost:
    count: 2
    host_groups:
      - ldap
```


Ansible Variables

Topics

- *Ansible Variables*
 - *Inputs*
 - *Built-ins*
 - *Defaults*

Inputs

The following variables can be set using ansible extra_vars, including in the [evars] section of linchpin.conf, to alter linchpin's default behavior.

topology

topology_file A set of rules, written in YAML, that define the way the provisioned systems should look after executing linchpin.

Generally, the *topology* and *topology_file* values are interchangeable, except after the file has been processed.

schema

schema_file JSON description of the format for the topology. (*schema_v3, schema_v4 are still available*)

layout

layout_file YAML definition for providing an ansible (currently) static inventory file, based upon the provided topology.

inventory

inventory_file If layout / layout_file is provided, this will be the location of the resulting ansible inventory.

linchpin_config if passed on the command line with `-c/--config`, should be an ini-style config file with linchpin default configurations (see BUILT-INS below for more information)

resources

resources_file File with the resource outputs in a JSON formatted file. Useful for teardown (destroy,down) actions depending on the provider.

workspace If provided, the above variables will be adjusted and mapped according to this value. Each path will use the following variables:

```
topology / topology_file = /<topologies_folder>
layout / layout_file = /<layouts_folder>
resources / resources_file = /resources_folder>
inventory / inventory_file = /<inventories_folder>

.. note:: schema is not affected by this pathing
```

If the `WORKSPACE` environment variable is set, it will be used here. If it is not, this variable can be set on the command line with `-w/--workspace`, and defaults to the location of the PinFile bring provisioned.

Built-ins

These variables **SHOULD NOT** be changed!

lp_path base path for linchpin playbooks and python api

lpconfig <lp_path>/linchpin.conf, unless overridden by *linchpin_config*

Defaults

While the variables here can also be passed as extra-vars, the values are the defaults and it is recommended not to change them. These values are defined in <lp_path>/linchpin.conf by default.

async (*boolean, default: False*)

Used to enable asynchronous provisioning/teardown

async_timeout (*int, default: 1000*)

How long the resource collection (formerly outputs_writer) process should wait

output (*boolean, default: True, previous: no_output*)

Controls whether resources will be written to the resources_file

check_mode (*boolean, default: no*)

This option does nothing at this time, though it may eventually be used for dry-run functionality based upon the provider

schemas_folder (*file_path, default: schemas*)

relative path to schemas

playbooks_folder (*file_path, default: provision*)

relative path to playbooks, only useful to the linchpin API and CLI

layouts_folder (*file_path, default: layouts*)

relative path to layouts

topologies_folder (*file_path, default: topologies*)

relative path to topologies

default_schemas_path (*file_path, default: <lp_path>/defaults/<schemas_folder>*)

default path to schemas, absolute path. Can be overridden by passing schema / schema_file.

default_playbooks_path (*file_path, default: <lp_path>/defaults/playbooks_folder*)

default path to playbooks location, only useful to the linchpin API and CLI

default_layouts_path (*file_path, default: <lp_path>/defaults/<layouts_folder>*)

default path to inventory layout files

default_topologies_path (*file_path, default: <lp_path>/defaults/<topologies_folder>*)

default path to topology files

default_resources_path (*file_path, default: <lp_path>/defaults/<resources_folder>, formerly: outputs*)

landing location for resources output data

default_inventories_path (*file_path, default: <lp_path>/defaults/<inventories_folder>*)

landing location for inventory outputs

See also:

[Glossary](#) Glossary

Python API Reference

This page contains the list of project's modules

linchpin module

The linchpin module contains calls to implement the Command Line Interface within linchpin. It uses the [Click](#) command line interface composer.

class `linchpin.LinchpinAliases` (*name=None, commands=None, **attrs*)

get_command (*ctx, name*)

Track aliases for specific commands and the commands and return the correct action.

list_commands (*ctx*)

Provide a list of available commands. Anything deprecated should not be listed

lp_aliases = {'down': 'destroy', 'drop': 'destroy', 'rise': 'up'}

lp_commands = ['init', 'up', 'destroy']

`linchpin.init()`

Initializes a linchpin project, which generates an example PinFile, and creates the necessary directory structure for topologies and layouts.

Parameters **ctx** – Context object defined by the `click.make_pass_decorator` method

`linchpin.up()`

Provisions nodes from the given target(s) in the given PinFile.

Parameters

- **ctx** – Context object defined by the `click.make_pass_decorator` method
- **pinfile** – path to pinfile (Default: `ctx.workspace`)
- **targets** – Provision ONLY the listed target(s). If omitted, ALL targets in the appropriate PinFile will be provisioned.

`linchpin.destroy()`

Destroys nodes from the given target(s) in the given PinFile.

Parameters

- **ctx** – Context object defined by the `click.make_pass_decorator` method
- **pinfile** – path to pinfile (Default: `ctx.workspace`)
- **targets** – Destroy ONLY the listed target(s). If omitted, ALL targets in the appropriate PinFile will be destroyed.

linchpin.api module

This page contains the list of project's modules

class `linchpin.api.LinchpinAPI` (*ctx*)

bind_to_hook_state (*callback*)

Function used by LinchpinHooksclass to add callbacks

Parameters **callback** – callback function

find_topology (*topology*)

Find the topology to be acted upon. This could be pulled from a registry.

Parameters **topology** – name of topology from PinFile to be loaded

get_cfg (*section=None, key=None, default=None*)

Get cfgs value(s) by section and/or key, or the whole cfgs object

Parameters

- **section** – section from ini-style config file
- **key** – key to get from config file, within section
- **default** – default value to return if nothing is found.

Does not apply if section is not provided.

get_evar (*key=None, default=None*)

Get the current evars (extra_vars)

Parameters

- **key** – key to use
- **default** – default value to return if nothing is found (default: None)

hook_state

getter function for hook_state property of the API object

lp_destroy (*pinfile, targets='all'*)

This function takes a list of targets, and performs a destructive teardown, including undefining nodes, according to the target.

See also:

lp_down - currently unimplemented

Parameters

- **pinfile** – Provided PinFile, with available targets,
- **targets** – A tuple of targets to destroy.

lp_down (*pinfile, targets='all'*)

This function takes a list of targets, and performs a shutdown on nodes in the target's topology. Only providers which support shutdown from their API (Ansible) will support this option.

CURRENTLY UNIMPLEMENTED

See also:

lp_destroy

Parameters

- **pinfile** – Provided PinFile, with available targets,
- **targets** – A tuple of targets to provision.

lp_drop (*pinfile, targets*)

DEPRECATED

An alias for lp_destroy. Used only for backward compatibility.

lp_rise (*pinfile, targets='all'*)

DEPRECATED

An alias for lp_up. Used only for backward compatibility.

lp_up (*pinfile, targets='all'*)

This function takes a list of targets, and provisions them according to their topology. If an layout argument is provided, an inventory will be generated for the provisioned nodes.

Parameters

- **pinfile** – Provided PinFile, with available targets,
- **targets** – A tuple of targets to provision.

run_playbook (*pinfile, targets='all', playbook='up'*)

This function takes a list of targets, and executes the given playbook (provision, destroy, etc.) for each provided target.

Parameters

- **pinfile** – Provided PinFile, with available targets,
- **targets** – A tuple of targets to run. (default: 'all')

set_cfg (*section, key, value*)

Set a value in cfgs. Does not persist into a file, only during the current execution.

Parameters

- **section** – section within ini-style config file
- **key** – key to use
- **value** – value to set into section within config file

set_evar (*key, value*)

Set a value into evars (extra_vars). Does not persist into a file, only during the current execution.

Parameters

- **key** – key to use
- **value** – value to set into evars

set_magic_vars ()

Function inbuilt to set magic vars for ansible context

linchpin.api.suppress_stdout (**args, **kws*)

This context manager provides tooling to make Ansible's Display class not output anything when used

linchpin.cli module

This page contains the list of project's modules

class `linchpin.cli.LinchpinCli (ctx)`

Bases: `linchpin.api.LinchpinAPI`

bind_to_hook_state (*callback*)

Function used by LinchpinHooksclass to add callbacks

Parameters **callback** – callback function

find_topology (*topology*)

Find the topology to be acted upon. This could be pulled from a registry.

Parameters **topology** – name of topology from PinFile to be loaded

get_cfg (*section=None, key=None, default=None*)

Get cfgs value(s) by section and/or key, or the whole cfgs object

Parameters

- **section** – section from ini-style config file
- **key** – key to get from config file, within section
- **default** – default value to return if nothing is found.

Does not apply if section is not provided.

get_evar (*key=None, default=None*)

Get the current evvars (extra_vars)

Parameters

- **key** – key to use
- **default** – default value to return if nothing is found (default: None)

hook_state

getter function for hook_state property of the API object

lp_destroy (*pinfile, targets='all'*)

This function takes a list of targets, and performs a destructive teardown, including undefining nodes, according to the target.

See also:

lp_down - currently unimplemented

Parameters

- **pinfile** – Provided PinFile, with available targets,
- **targets** – A tuple of targets to destroy.

lp_down (*pinfile, targets='all'*)

This function takes a list of targets, and performs a shutdown on nodes in the target's topology. Only providers which support shutdown from their API (Ansible) will support this option.

CURRENTLY UNIMPLEMENTED

See also:

lp_destroy

Parameters

- **pinfile** – Provided PinFile, with available targets,
- **targets** – A tuple of targets to provision.

lp_drop (*pinfile, targets*)

DEPRECATED

An alias for lp_destroy. Used only for backward compatibility.

lp_rise (*pinfile, targets='all'*)

DEPRECATED

An alias for lp_up. Used only for backward compatibility.

lp_up (*pinfile, targets='all'*)

This function takes a list of targets, and provisions them according to their topology. If an layout argument is provided, an inventory will be generated for the provisioned nodes.

Parameters

- **pinfile** – Provided PinFile, with available targets,
- **targets** – A tuple of targets to provision.

run_playbook (*pinfile, targets='all', playbook='up'*)

This function takes a list of targets, and executes the given playbook (provision, destroy, etc.) for each provided target.

Parameters

- **pinfile** – Provided PinFile, with available targets,
- **targets** – A tuple of targets to run. (default: 'all')

set_cfg (*section, key, value*)

Set a value in cfgs. Does not persist into a file, only during the current execution.

Parameters

- **section** – section within ini-style config file
- **key** – key to use
- **value** – value to set into section within config file

set_evar (*key, value*)

Set a value into evars (extra_vars). Does not persist into a file, only during the current execution.

Parameters

- **key** – key to use
- **value** – value to set into evars

set_magic_vars ()

Function inbuilt to set magic vars for ansible context

Glossary

The following is a list of terms used throughout the LinchPin documentation.

async (*boolean, default: False*)

Used to enable asynchronous provisioning/teardown

async_timeout (*int, default: 1000*)

How long the resource collection (formerly outputs_writer) process should wait

check_mode (*boolean, default: no*)

This option does nothing at this time, though it may eventually be used for dry-run functionality based upon the provider

default_schemas_path (*file_path, default: <lp_path>/defaults/<schemas_folder>*)

default path to schemas, absolute path. Can be overridden by passing schema / schema_file.

default_playbooks_path (*file_path, default: <lp_path>/defaults/playbooks_folder>*)

default path to playbooks location, only useful to the linchpin API and CLI

default_layouts_path (*file_path, default: <lp_path>/defaults/<layouts_folder>*)

default path to inventory layout files

default_topologies_path (*file_path, default: <lp_path>/defaults/<topologies_folder>*)

default path to topology files

default_resources_path (*file_path, default: <lp_path>/defaults/<resources_folder>, formerly: outputs*)

default landing location for resources output data

default_inventories_path (*file_path, default: <lp_path>/defaults/<inventories_folder>*)

default landing location for inventory outputs

hook Certain scripts can be called when a particular *hook* has been referenced in the *PinFile*. The currently available hooks are *preup*, *postup*, *predestroy*, and *postdestroy*.

inventory

inventory_file If layout / layout_file is provided, this will be the location of the resulting ansible inventory.

linchpin_config if passed on the command line with `-c/--config`, should be an ini-style config file with linchpin default configurations (see BUILT-INS below for more information)

layout

layout_file YAML definition for providing an ansible (currently) static inventory file, based upon the provided topology.

layouts_folder (*file_path, default: layouts*)

relative path to layouts

lp_path base path for linchpin playbooks and python api

lpconfig `<lp_path>/linchpin.conf`, unless overridden by *linchpin_config*

output (*boolean, default: True, previous: no_output*)

Controls whether resources will be written to the resources_file

PinFile A YAML file consisting of a *topology* and an optional *layout*, among other options. This file is used by the `linchpin` command-line, or Python API to determine what resources are needed for the current action.

playbooks_folder (*file_path, default: provision*)

relative path to playbooks, only useful to the linchpin API and CLI

provider A set of platform actions grouped together, which is provided by an external Ansible module. *openstack* would be a provider.

provision An action taken when resources are to be made available on a particular provider platform. Usually corresponds with the `linchpin up` command.

resources

resources_file File with the resource outputs in a JSON formatted file. Useful for teardown (destroy,down) actions depending on the provider.

schema JSON description of the format for the topology.

(schema_v3, schema_v4 are still available)

schemas_folder *(file_path, default: schemas)*

relative path to schemas

target Specified in the *PinFile*, the *target* references a *topology* and optional *layout* to be acted upon from the command-line utility, or Python API.

teardown An action taken when resources are to be made unavailable on a particular provider platform. Usually corresponds with the `linchpin destroy` command.

topologies_folder *(file_path, default: topologies)*

relative path to topologies

topology

topology_file A set of rules, written in YAML, that define the way the provisioned systems should look after executing `linchpin`.

Generally, the *topology* and *topology_file* values are interchangeable, except after the file has been processed.

topology_name Within a *topology_file*, the *topology_name* provides a way to identify the set of resources being acted upon.

workspace If provided, the above variables will be adjusted and mapped according to this value. Each path will use the following variables:

```
topology / topology_file = /<topologies_folder>
layout / layout_file = /<layouts_folder>
resources / resources_file = /resources_folder>
inventory / inventory_file = /<inventories_folder>
```

If the `WORKSPACE` environment variable is set, it will be used here. If it is not, this variable can be set on the command line with `-w/--workspace`, and defaults to the location of the `PinFile` being provisioned.

Note: schema is not affected by this pathing

See also:

[Ansible Variables](#)

[Source Code](#) LinchPin Source Code

[User Mailing List](#) Subscribe and participate. A great place for Q&A

[irc.freenode.net](#) #linchpin IRC chat channel

CHAPTER 2

Indices and tables

- `genindex`
- `modindex`
- `search`

I

`linchpin`, [15](#)
`linchpin.api`, [16](#)
`linchpin.cli`, [18](#)

A

async, [14](#), [20](#)
async_timeout, [14](#), [20](#)

B

bind_to_hook_state() (linchpin.api.LinchpinAPI
method), [16](#)
bind_to_hook_state() (linchpin.cli.LinchpinCli method),
[18](#)

C

check_mode, [14](#), [20](#)

D

default_inventories_path, [15](#), [20](#)
default_layouts_path, [14](#), [20](#)
default_playbooks_path, [14](#), [20](#)
default_resources_path, [14](#), [20](#)
default_schemas_path, [14](#), [20](#)
default_topologies_path, [14](#), [20](#)
destroy() (in module linchpin), [15](#)

F

find_topology() (linchpin.api.LinchpinAPI method), [16](#)
find_topology() (linchpin.cli.LinchpinCli method), [18](#)

G

get_cfg() (linchpin.api.LinchpinAPI method), [16](#)
get_cfg() (linchpin.cli.LinchpinCli method), [18](#)
get_command() (linchpin.LinchpinAliases method), [15](#)
get_evar() (linchpin.api.LinchpinAPI method), [16](#)
get_evar() (linchpin.cli.LinchpinCli method), [18](#)

H

hook, [20](#)
hook_state (linchpin.api.LinchpinAPI attribute), [16](#)
hook_state (linchpin.cli.LinchpinCli attribute), [18](#)

I

init() (in module linchpin), [15](#)
inventory, [13](#), [20](#)
inventory_file, [13](#), [20](#)

L

layout, [13](#), [20](#)
layout_file, [13](#), [20](#)
layouts_folder, [14](#), [20](#)
linchpin (module), [15](#)
linchpin.api (module), [16](#)
linchpin.cli (module), [18](#)
linchpin_config, [13](#), [20](#)
LinchpinAliases (class in linchpin), [15](#)
LinchpinAPI (class in linchpin.api), [16](#)
LinchpinCli (class in linchpin.cli), [18](#)
list_commands() (linchpin.LinchpinAliases method), [15](#)
lp_aliases (linchpin.LinchpinAliases attribute), [15](#)
lp_commands (linchpin.LinchpinAliases attribute), [15](#)
lp_destroy() (linchpin.api.LinchpinAPI method), [16](#)
lp_destroy() (linchpin.cli.LinchpinCli method), [18](#)
lp_down() (linchpin.api.LinchpinAPI method), [16](#)
lp_down() (linchpin.cli.LinchpinCli method), [18](#)
lp_drop() (linchpin.api.LinchpinAPI method), [17](#)
lp_drop() (linchpin.cli.LinchpinCli method), [19](#)
lp_path, [14](#), [20](#)
lp_rise() (linchpin.api.LinchpinAPI method), [17](#)
lp_rise() (linchpin.cli.LinchpinCli method), [19](#)
lp_up() (linchpin.api.LinchpinAPI method), [17](#)
lp_up() (linchpin.cli.LinchpinCli method), [19](#)
lpconfig, [14](#), [20](#)

O

output, [14](#), [20](#)

P

PinFile, [20](#)
playbooks_folder, [14](#), [20](#)
provider, [21](#)

provision, [21](#)

R

resources, [13](#), [21](#)

resources_file, [13](#), [21](#)

run_playbook() (linchpin.api.LinchpinAPI method), [17](#)

run_playbook() (linchpin.cli.LinchpinCli method), [19](#)

S

schema, [13](#), [21](#)

schema_file, [13](#)

schemas_folder, [14](#), [21](#)

set_cfg() (linchpin.api.LinchpinAPI method), [17](#)

set_cfg() (linchpin.cli.LinchpinCli method), [19](#)

set_evar() (linchpin.api.LinchpinAPI method), [17](#)

set_evar() (linchpin.cli.LinchpinCli method), [19](#)

set_magic_vars() (linchpin.api.LinchpinAPI method), [17](#)

set_magic_vars() (linchpin.cli.LinchpinCli method), [19](#)

suppress_stdout() (in module linchpin.api), [17](#)

T

target, [21](#)

teardown, [21](#)

topologies_folder, [14](#), [21](#)

topology, [13](#), [21](#)

topology_file, [13](#), [21](#)

topology_name, [21](#)

U

up() (in module linchpin), [15](#)

W

workspace, [13](#), [21](#)