

---

# **elasticcluster Documentation**

***Release 1.0***

**Grid Computing Competence Centre, University of Zurich**

April 15, 2016



<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Features</b>	<b>3</b>
<b>3</b>	<b>Architecture</b>	<b>5</b>
<b>4</b>	<b>Table of Contents</b>	<b>7</b>
4.1	Installation . . . . .	7
4.2	Configuration . . . . .	9
4.3	Usage . . . . .	17
4.4	Playbooks distributed with elasticcluster . . . . .	25
4.5	Elasticcluster programming API . . . . .	30
<b>5</b>	<b>Indices and tables</b>	<b>47</b>
	<b>Python Module Index</b>	<b>49</b>



---

# Introduction

---

*elasticcluster* aims to provide a user-friendly command line tool to create, manage and setup computational clusters hosted on cloud infrastructures (like [Amazon's Elastic Compute Cloud EC2](#), [Google Compute Engine](#) or an [OpenStack](#) cloud). Its main goal is to get your own private cluster up and running with just a few commands; a [YouTube video](#) demoes the basic features of *elasticcluster*.

This project is an effort of the [Grid Computing Competence Center](#) at the [University of Zurich](#), licensed under the [GNU General Public License version 3](#).



---

## Features

---

*elasticcluster* is in active development, but the following features at the current state:

- Simple configuration file to define cluster templates
- Can start and manage multiple independent clusters at the same time
- **Automated cluster setup:**
  - use [Debian GNU/Linux](#), [Ubuntu](#), or [CentOS](#) as a base operating system
  - supports multiple batch systems, including [SLURM](#), [Grid Engine](#) or [Torque/PBS](#)
  - supports [Hadoop](#) clusters
  - add useful tools like [Ganglia](#) for monitoring...
  - ...or anything that you can install with an [Ansible](#) playbook!
- Grow a running cluster

*elasticcluster* is currently in active development: please use the GitHub issue tracker to [file enhancement requests](#) and [ideas](#)





---

# Architecture

---

The architecture of elasticcluster is quite simple: the configuration file in `~/elasticcluster/config` defines a set of *cluster configurations* and information on how to access a specific cloud webservice (including access id and secret keys).

Using the command line (or, very soon, a simple API), you can start a cluster and override some of the default values, like the number of nodes you want to fire up. Elasticcluster will use the [boto library](#) to connect to the desired cloud, start the virtual machines and wait until they are accessible via ssh.

After all the virtual machines are up and running, elasticcluster will use [ansible](#) to configure them.

If you do a *resize* of the cluster (currently only growing a cluster is fully supported) new virtual machines will be created and again [ansible](#) will run for *all* of the virtual machines, in order to properly add the new hosts to the cluster.



---

## Table of Contents

---

### 4.1 Installation

*elasticsearch* is a [Python](#) program; Python version 2.6 is required to run it.

The easiest way to install *elasticsearch* is using [pip](#), this will install the latest **stable** release from the [PyPI](#) website. The following section: *Installing from PyPI* will explain you how to do it.

If you instead want to test the *development* version, go to the *Installing from github* section.

In both cases, it's strongly suggested to install *elasticsearch* in a [python virtualenv](#), so that you can easily uninstall or upgrade *elasticsearch*.

#### 4.1.1 Installing from PyPI

It's quite easy to install *elasticsearch* using [pip](#); the command below is all you need to install *elasticsearch* on your system:

```
pip install elasticsearch
```

If you want to run *elasticsearch* from source you have to **install Ansible first**:

```
pip install ansible
python setup.py install
```

#### 4.1.2 Installing from github

The source code of *elasticsearch* is [github](#), if you want to test the latest development version you can clone the [github elasticsearch repository](#).

You need the `git` command in order to be able to clone it, and we suggest you to use [python virtualenv](#) in order to create a controlled environment in which you can install *elasticsearch* as normal user.

Assuming you already have `virtualenv` installed on your machine, you first need to create a `virtualenv` and install *ansible*, which is needed by *elasticsearch*:

```
virtualenv elasticsearch
. elasticsearch/bin/activate
pip install ansible
```

Then you have to download the software. We suggest you to download it *within* the created `virtualenv`:

```
cd elasticcluster
git clone git://github.com/gc3-uzh-ch/elasticcluster.git src
cd src
python setup.py install
```

Now the `elasticcluster` should be available in your current environment.

### 4.1.3 Notes on MacOSX installation

When installing *elasticcluster* on MacOSX operating systems you may get some errors while running *python setup.py install*, because *pip* is not always able to automatically resolve the dependencies.

In these cases, you need to find the package that is failing and install it manually using *pip*.

For instance, if during the installation you get something like:

```
Running requests-2.4.3/setup.py -q bdist_egg --dist-dir /var/folders/tZ/tZ2B3RaeGVq7+ptdJIbdj+++TI/
Adding requests 2.4.3 to easy-install.pth file
Traceback (most recent call last):
  File "setup.py", line 109, in <module>
    'elasticcluster = elasticcluster.main:main',
  File "/System/Library/Frameworks/Python.framework/Versions/2.6/lib/python2.6/distutils/core.py", line 151, in
    dist.run_commands()
  File "/System/Library/Frameworks/Python.framework/Versions/2.6/lib/python2.6/distutils/dist.py", line 954, in
    self.run_command(cmd)
  File "/System/Library/Frameworks/Python.framework/Versions/2.6/lib/python2.6/distutils/dist.py", line 973, in
    cmd_obj.run()
  File "/Users/michela/elasticcluster/build/setuptools/setuptools/command/install.py", line 65, in run
  File "/Users/michela/elasticcluster/build/setuptools/setuptools/command/install.py", line 115, in do_egg_install
  File "/Users/michela/elasticcluster/build/setuptools/setuptools/command/easy_install.py", line 360, in
    find_packages

  File "/Users/michela/elasticcluster/build/setuptools/setuptools/command/easy_install.py", line 576, in
    find_packages

  File "/Users/michela/elasticcluster/build/setuptools/setuptools/command/easy_install.py", line 627, in
    find_packages

  File "/Users/michela/elasticcluster/build/setuptools/setuptools/command/easy_install.py", line 682, in
    find_packages

  File "/Users/michela/elasticcluster/build/setuptools/pkg_resources.py", line 631, in resolve
    dist = best[req.key] = env.best_match(req, ws, installer)
  File "/Users/michela/elasticcluster/build/setuptools/pkg_resources.py", line 871, in best_match
    return self.obtain(req, installer)
  File "/Users/michela/elasticcluster/build/setuptools/pkg_resources.py", line 883, in obtain
    return installer(requirement)
  File "/Users/michela/elasticcluster/build/setuptools/setuptools/command/easy_install.py", line 595, in
    fetch_build_egg

  File "/Users/michela/elasticcluster/build/setuptools/setuptools/command/easy_install.py", line 627, in
    fetch_build_egg

  File "/Users/michela/elasticcluster/build/setuptools/setuptools/command/easy_install.py", line 659, in
    fetch_build_egg

  File "/Users/michela/elasticcluster/build/setuptools/setuptools/command/easy_install.py", line 532, in
    fetch_build_egg

  File "/Users/michela/elasticcluster/build/setuptools/setuptools/command/easy_install.py", line 734, in
    fetch_build_egg

  File "/private/var/folders/tZ/tZ2B3RaeGVq7+ptdJIbdj+++TI/-Tmp-/easy_install-qch0dG/python-keystoneclient-0.9.0
AttributeError: 'NoneType' object has no attribute 'get_script_header'
```

you probably need to install *pbr* manually using:

```
pip install pbr
```

In some MacOSX version, even if the installation *seems* to succeed, you may get the following error the first time you run *elasticcluster*:

```
Traceback (most recent call last):
  File "/Users/michela/el2/bin/elasticcluster", line 9, in <module>
    load_entry_point('elasticcluster==1.1-dev', 'console_scripts', 'elasticcluster')()
  File "/Users/michela/el2/build/setuptools/pkg_resources.py", line 356, in load_entry_point
    return get_distribution(dist).load_entry_point(group, name)
  File "/Users/michela/el2/build/setuptools/pkg_resources.py", line 2431, in load_entry_point
    return ep.load()
  File "/Users/michela/el2/build/setuptools/pkg_resources.py", line 2147, in load
    [ '__name__' ])
  File "build/bdist.macosx-10.6-universal/egg/elasticcluster/__init__.py", line 33, in <module>
  File "build/bdist.macosx-10.6-universal/egg/elasticcluster/providers/gce.py", line 37, in <module>
  File "build/bdist.macosx-10.6-universal/egg/apiclient/discovery.py", line 52, in <module>
  File "build/bdist.macosx-10.6-universal/egg/apiclient/errors.py", line 27, in <module>
ImportError: No module named anyjson
```

In this case, the issue is caused by *google-api-python-client*, and you should: 1) uninstall it using *pip uninstall* 2) reinstall it using *pip install* 3) re-run *elasticcluster* installation

```
pip uninstall google-api-python-client
[...]
pip install google-api-python-client
[...]
python setup.py install
```

## 4.2 Configuration

All the information about how to access a cloud provider and how to setup a cluster is stored in a configuration file. The default configuration file is stored in your home directory: `~/.elasticcluster/config` but you can specify a different location from the command line with the `-c` option.

When *elasticcluster* is run for the first time, if no configuration file is found it will copy a [template configuration file](#) in `~/.elasticcluster/config`. Such template is fully commented and self documented.

Therefore, after installing *elasticcluster* for the first time, we suggest to run the following command:

```
elasticcluster list-templates
```

If you don't already have a configuration file, this command will create one for you. Of course, the file is not complete, as it does not contain any authentication information, and you will get an error similar to the following:

```
WARNING:gc3.elasticcluster:Deploying default configuration file to /home/antonio/.elasticcluster/config
WARNING:gc3.elasticcluster:Ignoring Cluster `ipython`: required key not provided @ data['image_user']
WARNING:gc3.elasticcluster:Ignoring cluster `ipython`.
Error validating configuration file '/home/antonio/.elasticcluster/config': `required key not provided
```

In this case, you have to edit the configuration file in `~/.elasticcluster/config` and update it with the correct values.

Please refer to the following section to understand the syntax of the configuration file and to know which options you need to set in order to use *elasticcluster*.

### 4.2.1 Basic syntax of the configuration file

The file is parsed by ConfigParser module and has a syntax similar to Microsoft Windows INI files.

It consists of *sections* led by a `[sectiontype/name]` header and followed by lines in the form:

```
key=value
```

Section names are in the form `[type/name]` wher *type* must be one of:

**cloud** define a cloud provider

**login** define a way to access a virtual machine

**setup** define a way to setup the cluster

**cluster** define the composition of a cluster. It contains references to the other sections.

**cluster/<clustername>** override configuration for specific group of nodes within a cluster

You must define at least one for each section types in order to have a valid configuration file.

### 4.2.2 Cloud Section

A cloud section named `<name>` starts with:

```
[cloud/<name>]
```

The cloud section defines all properties needed to connect to a specific cloud provider.

You can define as many cloud sections you want, assuming you have access to different cloud providers and want to deploy different clusters in different clouds. The mapping between cluster and cloud provider is done in the *cluster* section (see later).

Currently two cloud providers are available: - boto: supports OpenStack and Amazon EC2 - google: supports Google Compute Engine Therefore the following configuration option needs to be set in the cloud section:

provider

the driver to use to connect to the cloud provider. *ec2\_boto*, *openstack* or *google*

#### Valid configuration keys for *boto*

ec2\_url

the url of the EC2 endpoint. For Amazon is probably something like:

```
https://ec2.us-east-1.amazonaws.com
```

replace `us-east-1` with the zone you want to use while for OpenStack you can get it from the web interface

ec2\_access\_key

the access key (also known as access id) your cloud provider gave you to access its cloud resources.

ec2\_secret\_key

the secret key (also known as secret id) your cloud provider gave you to access its cloud resources.

ec2\_region

the availability zone you want to use.

`request_floating_ip`

request assignment of a floating IP when the instance is started. Valid values: *True*, *False*. Some cloud providers does not automatically assign a public IP to the instances, but this is often needed if you want to connect to the VM from outside. Setting `request_floating_ip` to *True* will force *elasticsearch* to request such a floating IP if the instance doesn't get one automatically.

### Valid configuration keys for *google*

`gce_client_id`

The API client id generated in the Google API Console

`gce_client_secret`

The API client secret generated in the Google API Console

`gce_project_id`

The project id of your Google Compute Engine project

### Valid configuration keys for *openstack*

`auth_url:`

The URL of the keystone service (main entry point for OpenStack clouds)

`username`

OpenStack username

`password`

OpenStack password

`project_name`

OpenStack project to use (also known as *tenant*)

`region_name`

OpenStack region (optional)

`request_floating_ip`

request assignment of a floating IP when the instance is started. Valid values: *True*, *False*. Some cloud providers does not automatically assign a public IP to the instances, but this is often needed if you want to connect to the VM from outside. Setting `request_floating_ip` to *True* will force *elasticsearch* to request such a floating IP if the instance doesn't get one automatically.

### Examples

For instance, to connect to the [Hobbes private cloud](#) of the [University of Zurich](#) you can use the following:

```
[cloud/hobbes]
provider=ec2_boto
ec2_url=http://hobbes.gc3.uzh.ch:8773/services/Cloud
ec2_access_key=****REPLACE WITH YOUR ACCESS ID****
ec2_secret_key=****REPLACE WITH YOUR SECRET KEY****
ec2_region=nova
auto_ip_assignment=True
```

For Amazon instead (region us-east-1) you can use:

```
[cloud/amazon-us-east-1]
provider=ec2_boto
ec2_url=https://ec2.us-east-1.amazonaws.com
ec2_access_key=****REPLACE WITH YOUR ACCESS ID****
ec2_secret_key=****REPLACE WITH YOUR SECRET KEY****
ec2_region=us-east-1
```

For Google Compute Engine you can use:

```
[cloud/google]
provider=google
gce_client_id=****REPLACE WITH YOUR CLIENT ID****
gce_client_secret=****REPLACE WITH YOUR SECRET KEY****
gce_project_id=****REPLACE WITH YOUR PROJECT ID****
```

## OpenStack users

From the horizon web interface you can download a file containing your EC2 credentials by logging into your provider web interface and clicking on:

“*settings*“

=> “*EC2 Credentials*“ => “*Download EC2 Credentials*“

The `ec2rc.sh` file will contain some values. Update the configuration file:

- `ec2_url` using the value of the variable `EC2_URL`
- `ec2_access_key` using the value of the variable `EC2_ACCESS_KEY`
- `ec2_secret_key` using the value of the variable `EC2_SECRET_KEY`

## Google Compute Engine users

To generate a `client_id` and `client_secret` to access the Google Compute Engine visit the following page:  
<https://code.google.com/apis/console/>

1. Select the project defined as `gce_project_id`
2. Navigate to *API Access*
3. Click create another client id
4. Select *Installed Application -> Other*
5. After clicking the *Create* button you'll see your `client_id` and `secret_key` in the list of available client ids

**Please note:** you have to set your google username in the login section below as `image_user` to be able to use Google Compute Engine

### 4.2.3 Login Section

A login section named `<name>` starts with:

```
[login/<name>]
```



This section contains information on how to access the instances started on the cloud, including the user and the SSH keys to use.

Some of the values depend on the image you specified in the *cluster* section. Values defined here also can affect the *setup* section and the way the system is setup.

### Mandatory configuration keys

`image_user`

the remote user you must use to connect to the virtual machine. In case you're using Google Compute Engine you have to set your username here. So if your gmail address is [karl.marx@gmail.com](mailto:karl.marx@gmail.com), your username is karl.marx

`image_sudo`

Can be *True* or *False*. *True* means that on the remote machine you can execute commands as root by running the *sudo* program.

`image_user_sudo`

the login name of the administrator. Use *root* unless you know what you are doing...

`user_key_name`

name of the *keypair* to use on the cloud provider. If the keypair does not exist it will be created by elasticcluster.

`user_key_private`

file containing a valid RSA or DSA private key to be used to connect to the remote machine. Please note that this must match the `user_key_public` file (RSA and DSA keys go in pairs). Also note that Amazon does not accept DSA keys but only RSA ones.

`user_key_public`

file containing the RSA/DSA public key corresponding to the `user_key_private` private key. See `user_key_private` for more details.

### Examples

For a typical Ubuntu machine, both on Amazon and most OpenStack providers, these values should be fine:

```
[login/ubuntu]
image_user=ubuntu
image_user_sudo=root
image_sudo=True
user_key_name=elasticcluster
user_key_private=~/.ssh/id_rsa
user_key_public=~/.ssh/id_rsa.pub
```

while for Hobbes appliances you will need to use the *gc3-user* instead:

```
[login/gc3-user]
image_user=gc3-user
image_user_sudo=root
image_sudo=True
user_key_name=elasticcluster
user_key_private=~/.ssh/id_rsa
user_key_public=~/.ssh/id_rsa.pub
```

## 4.2.4 Setup Section

A setup section named `<name>` starts with:

```
[setup/<name>]
```

This section contain information on *how to setup* a cluster. After the cluster is started, elasticcluster will run a `setup` provider in order to configure it.

### Mandatory configuration keys

`provider`

the type of setup provider. So far, only *ansible* is supported.

### Ansible-specific mandatory configuration keys

The following configuration keys are only valid if *provider* is *ansible*.

`<class>_groups`

Comma separated list of ansible groups the specific `<class>` will belong to. For each `<class>_nodes` in a `[cluster/]` section there should be a `<class>_groups` option to configure that specific class of nodes with the ansible groups specified.

If you are setting up a standard HPC cluster you probably want to have only two main groups: *frontend\_groups* and *compute\_groups*.

To configure a slurm cluster, for instance, you have the following available groups:

**slurm\_master** configure this machine as slurm masternode

**slurm\_clients** compute nodes of a slurm cluster

**ganglia\_master** configure as ganglia web frontend. On the master, you probably want to define *ganglia\_monitor* as well

**ganglia\_monitor** configure as ganglia monitor.

You can combine more groups together, but of course not all combinations make sense. A common setup is, for instance:

```
frontend_groups=slurm_master,ganglia_master,ganglia_monitor
compute_groups=slurm_clients,ganglia_monitor
```

This will configure the frontend node as slurm master and ganglia frontend, and the compute nodes as clients for both slurm and ganglia frontend.

A full list of the available groups is available at the [Playbooks distributed with elasticcluster](#) page.

`<class>_var_<varname>`

an entry of this type will define a variable called `<varname>` for the specific `<class>` and add it to the ansible inventory file.

`playbook_path`

Path to the playbook to use when configuring the system. The default value printed here points to the playbook distributed with elasticcluster. The default value points to the playbooks distributed with elasticcluster.

## Examples

Some (working) examples:

```
[setup/ansible-slurm]
provider=ansible
frontend_groups=slurm_master
compute_groups=slurm_clients

[setup/ansible-gridengine]
provider=ansible
frontend_groups=gridengine_master
compute_groups=gridengine_clients

[setup/ansible-pbs]
provider=ansible
frontend_groups=pbs_master,maui_master
compute_groups=pbs_clients

[setup/ansible_matlab]
# Please note that this setup assumes you already have matlab
# installed on the image that is being used.
provider=ansible
frontend_groups=mdce_master,mdce_worker,ganglia_monitor,ganglia_master
worker_groups=mdce_worker,ganglia_monitor
```

### 4.2.5 Cluster Section

A cluster section named <name> starts with:

```
[cluster/<name>]
```

The cluster section defines a *template* for a cluster. This section has references to each one of the other sections and define the image to use, the default number of compute nodes and the security group.

#### Mandatory configuration keys

cloud

the name of a valid *cloud* section. For instance *hobbes* or *amazon-us-east-1*

login

the name of a valid *login* section. For instance *ubuntu* or *gc3-user*

setup\_provider

the name of a valid *setup* section. For instance, *ansible-slurm* or *ansible-pbs*

image\_id

image id in *ami* format. If you are using OpenStack, you need to run *euca-describe-images* to get a valid *ami-\** id.

flavor

the image type to use. Different cloud providers call it differently, could be *instance type*, *instance size* or *flavor*.

security\_group

Security group to use when starting the instance.

`<class>_nodes`

the number of nodes of type `<class>`. These configuration options will define the composition of your cluster. A very common configuration will include only two group of nodes:

**frontend\_nodes** the queue manager and frontend of the cluster. You probably want only one.

**compute\_nodes** the worker nodes of the cluster.

Each `<class>_nodes` group is configured using the corresponding `<class>_groups` configuration option in the `[setup/. . .]` section.

`ssh_to`

*ssh* and *sftp* nodes will connect to only one node. This is the first of the group specified in this configuration option, or the first node of the first group in alphabetical order. For instance, if you don't set any value for *ssh\_to* and you defined two groups: *frontend\_nodes* and *compute\_nodes*, the *ssh* and *sftp* command will connect to *compute001* which is the first *compute\_nodes* node. If you specify *frontend*, instead, it will connect to *frontend001* (or the first node of the *frontend* group).

## Optional configuration keys

`image_userdata`

shell script to be executed (as root) when the machine starts. This is usually not needed because the *ansible* provider works on *vanilla* images, but if you are using other setup providers you may need to execute some command to bootstrap it.

`network_ids`

comma separated list of network IDs the nodes of the cluster will be connected to. Only supported when the cloud provider is *openstack*

## Examples

Some (working) examples:

```
[cluster/slurm]
cloud=hobbes
login=gc3-user
setup_provider=ansible-slurm
security_group=default
# Ubuntu image
image_id=ami-00000048
flavor=m1.small
frontend_nodes=1
compute_nodes=2
frontend_class=frontend

[cluster/torque]
cloud=hobbes
frontend_nodes=1
compute_nodes=2
frontend_class=frontend
security_group=default
# CentOS image
image_id=ami-0000004f
```

```

flavor=m1.small
login=gc3-user
setup_provider=ansible-pbs

[cluster/aws-slurm]
cloud=amazon-us-east-1
login=ubuntu
setup_provider=ansible-slurm
security_group=default
# ubuntu image
image_id=ami-90a21cf9
flavor=m1.small
frontend=1
compute=2

[cluster/matlab]
cloud=hobbes
setup_provider=ansible_matlab
security_group=default
image_id=ami-00000099
flavor=m1.medium
frontend_nodes=1
worker_nodes=10
image_userdata=
ssh_to=frontend

```

#### 4.2.6 Cluster node section

A *cluster node* for the node type <nodetype> of the cluster <name> starts with:

```
[cluster/<name>/<nodetype>]
```

This section allows you to override some configuration values for specific group of nodes. Assume you have a standard slurm cluster with a frontend which is used as manager node and nfs server for the home directories, and a set of compute nodes.

You may want to use different flavors for the frontend and the compute nodes, since for the first you need more space and you don't need many cores or much memory, while the compute nodes may requires more memory and more cores but are not eager about disk space.

This is achieved defining, for instance, a *bigdisk* flavor (the name is just fictional) for the frontend and *8cpu32g* for the compute nodes. Your configuration will thus look like:

```

[cluster/slurm]
...
flavor=8cpu32g
frontend_nodes=1
compute_nodes=10

[cluster/slurm/frontend]
flavor=bigdisk

```

### 4.3 Usage

The syntax of the elasticcluster command is:

```
elasticcluster [-v] [-s PATH] [-c PATH] [subcommand] [subcommand args and opts]
```

The following options are general and are accepted by any subcommand:

**-h, --help** Show an help message and exits.

**-v, --verbose** Adding one or more *-v* will increase the verbosity. Usually elasticcluster creates new VMs in parallel, to speedup the process, but if you run it with at least *four* *-v* options, elasticcluster will not fork and will start the VMs sequentially. Useful for debugging.

**-s PATH, --storage PATH**

Path to the storage folder. This directory is used to store information about the cluster which are running. By default this is “`~/elasticcluster/storage`”

**WARNING:** If you delete this directory elasticcluster will not be able to access the cluster anymore!

**-c PATH, --config PATH**

Path to the configuration file. By default this is `~/elasticcluster/config`

elasticcluster provides multiple *subcommands* to start, stop, resize, inspect your clusters. The available subcommands are:

**start** Create a cluster using one of the configured cluster tmlate.

**stop** Stop a cluster and all associated VM instances.

**list** List all clusters that are currently running.

**list-nodes** Show information about the nodes in a specific started cluster.

**list-templates** Show the available cluster configurations, as defined in the configuration file.

**setup** Run ansible to configure the cluster.

**resize** Resize a cluster by adding or removing nodes.

**ssh** Connect to the frontend of the cluster using the *ssh* command.

**sftp** Open an SFTP session to the cluster frontend host.

An help message explaining the available options and subcommand of *elasticcluster* is available by running:

```
elasticcluster -h
```

Options and arguments accepted by a specific subcommand *<cmd>* is available by running:

```
elasticcluster <cmd> -h
```

### 4.3.1 The start command

This command will start a new cluster using a specific cluster configuration, defined in the configuration file. You can start as many clusters you want using the same cluster configuration, by providing different *--name* options.

Basic usage of the command is:

```
usage: elasticcluster start [-h] [-v] [-n CLUSTER_NAME]
                           [--nodes N1:GROUP[,N2:GROUP2,...]] [--no-setup]
                           cluster
```

*cluster* is the name of a *cluster* section in the configuration file. For instance, to start the cluster defined by the section `[cluster/slurm]` you must run the command:

```
elasticcluster start slurm
```

The following options are available:

**-h, --help** Show an help message and exits.

**-v, --verbose** Adding one or more *-v* will increase the verbosity accordingly.

**-n CLUSTER\_NAME, --name CLUSTER\_NAME** Name of the cluster. By default this is the same as the cluster configuration name.

**--nodes N1:GROUP [, N2:GROUP2, ...]**

This option allow you to override the values stored in the configuration file, by starting a different number of hosts fore each group.

Assuming you defined, for instance, a cluster with the following type of nodes in the configuration file:

```
hadoop-data_nodes=4
hadoop-task_nodes=4
```

and you want to run instead 10 data nodes and 10 task nodes, you can run elasticcluster with option:

```
elasticcluster ... --nodes 10:hadoop-data,10:hadoop-task
```

**--no-setup** By default elasticcluster will automatically run the **setup** command after all the virtual machines are up and running. This option prevent the *setup* step to be run and will leave the cluster unconfigured.

When you start a new cluster, elasticcluster will:

- create the requested/configured number of virtual machines.
- wait until *all* the virtual machines are started.
- wait until *elasticcluster* is able to connect to *all* the virtual machines using *ssh*.
- run ansible on all the virtual machines (unless *--no-setup* option is given).

This process can take several minutes, depending on the load of the cloud, the configuration of the cluster and your connection speed. *Elasticcluster* usually print very few information on what's happening, if you run it with *-v* it will display a more verbose output (including output of ansible command) to help you understanding what is actually happening.

After the setup process is done a summary of the created cluster is printed, similar to the following:

```
Cluster name:      slurm
Cluster template:  slurm
Frontend node:  frontend001
- compute nodes:  2
- frontend nodes: 1
```

To login on the frontend node, run the command:

```
elasticcluster ssh slurm
```

To upload or download files to the cluster, use the command:

```
elasticcluster sftp slurm
```

The first line tells you the name of the cluster, which is the one you are supposed to use with the **stop**, **list-nodes**, **resize**, **ssh** and **sftp** commands.

The second line specifies the cluster configuration section used to configure the cluster (in this case, for instance, the section `[cluster/slurm]` has been used)

The `Frontend` node line shows which node is used for the `ssh` and `sftp` commands, when connecting to the cluster.

Then a list of how many nodes of each type have been started

The remaining lines describe how to connect to the cluster either by opening an interactive shell to run commands on it, or an `sftp` session to upload and download files.

### 4.3.2 The `stop` command

The **`stop`** command will terminate all the instances running and delete all information related to the cluster saved on the local disk.

**WARNING:** elasticcluster doesn't do any kind of test to check if the cluster is *used*!

Basic usage of the command is:

```
usage: elasticcluster stop [-h] [-v] [--force] [--yes] cluster
```

Like the **`start`** command, `cluster` is the name of a *cluster* section in the configuration file.

The following options are available:

**`-h, --help`** Show an help message and exits.

**`-v, --verbose`** Adding one or more `-v` will increase the verbosity accordingly.

**`--force`**

If some of the virtual machines fail to terminate (for instance because they have been terminated already not by elasticcluster), this command will ignore these errors and will force termination of all the other instances.

**`--yes`**

Since stopping a cluster is a possibly disruptive action, elasticcluster will always ask for confirmation before doing any modification, unless this option is given.

### 4.3.3 The `list` command

The **`list`** command print a list of all the cluster that have been started. For each cluster, it will print a few information including the cloud used and the number of nodes started for each node type:

```
$ elasticcluster list

The following clusters have been started.
Please note that there's no guarantee that they are fully configured:

centossge
-----
name:          centossge
template:      centossge
cloud:         hobbes
- frontend nodes: 1
- compute nodes: 2

slurm
-----
name:          slurm
template:      slurm
cloud:         hobbes
```



```

- frontend nodes: 1
- compute nodes: 2

slurm13.04
-----
name:          slurm13.04
template:      slurm13.04
cloud:         hobbles
- frontend nodes: 1
- compute nodes: 2

```

### 4.3.4 The `list-nodes` command

The **list-nodes** command print information on the nodes belonging to a specific cluster.

Basic usage of the command is:

```
usage: elasticcluster list-nodes [-h] [-v] [-u] cluster
```

`cluster` is the name of a cluster that has been *started* previously.

The following options are available:

**-h, --help** Show an help message and exits.

**-v, --verbose** Adding one or more `-v` will increase the verbosity accordingly.

**-u, --update**

By default `elasticcluster list-nodes` will not contact the EC2 provider to get up-to-date information, unless `-u` option is given.

Example:

```

$ elasticcluster list-nodes centossge

Cluster name:      centossge
Cluster template:  centossge
Frontend node:     frontend001
- frontend nodes:  1
- compute nodes:   2

To login on the frontend node, run the command:

    elasticcluster ssh centossge

To upload or download files to the cluster, use the command:

    elasticcluster sftp centossge

frontend nodes:

- frontend001
  public IP:    130.60.24.61
  private IP:   10.10.10.36
  instance id:  i-0000299f
  instance flavor: m1.small

compute nodes:

```

```
- compute001
  public IP: 130.60.24.44
  private IP: 10.10.10.17
  instance id: i-0000299d
  instance flavor: m1.small

- compute002
  public IP: 130.60.24.48
  private IP: 10.10.10.29
  instance id: i-0000299e
  instance flavor: m1.small
```

### 4.3.5 The `list-templates` command

The **list-templates** command print a list of all the available templates defined in the configuration file with a few information for each one of them.

Basic usage of the command is:

```
usage: elasticcluster list-templates [-h] [-v] [clusters [clusters ...]]
```

*clusters* is used to limit the clusters to be listed and uses a globbing-like pattern matching. For instance, to show all the cluster templates that contains the word `slurm` in their name you can run the following:

```
$ elasticcluster list-templates *slurm*
11 cluster templates found.

name:      aws-slurm
cloud:     aws
compute nodes: 2
frontend nodes: 1

name:      slurm
cloud:     hobbes
compute nodes: 2
frontend nodes: 1

name:      slurm_xl
cloud:     hobbes
compute nodes: 2
frontend nodes: 1

name:      slurm13.04
cloud:     hobbes
compute nodes: 2
frontend nodes: 1
```

### 4.3.6 The `setup` command

The **setup** command will run *ansible* on the desired cluster once again. It is usually needed only when you customize and update your playbooks, in order to re-configure the cluster, since the **start** command already run *ansible* when all the machines are started.

Basic usage of the command is:

```
usage: elasticsearch setup [-h] [-v] cluster
```

`cluster` is the name of a cluster that has been *started* previously.

The following options are available:

**-h, --help** Show an help message and exits.

**-v, --verbose** Adding one or more `-v` will increase the verbosity accordingly.

### 4.3.7 The `resize` command

The **resize** command allow you to add or remove nodes from a started cluster. Please, be warned that **this feature is still experimental**, and while adding nodes is usually safe, removing nodes can be disruptive and can leave the cluster in an unknown state.

Moreover, there is currently no way to decide *which nodes* can be removed from a cluster, therefore if you shrink a cluster **you must ensure** that any node of that type can be removed safely and no job is running on it.

When adding nodes, you have to specify the *type* of the node and the number of node you want to add. Then, elasticsearch will basically re-run the *start* and *setup* steps:

- create the requested/configured number of virtual machines.
- wait until *all* the virtual machines are started.
- wait until *elasticsearch* is able to connect to *all* the virtual machines using *ssh*.
- run ansible on all the virtual machines, including the virtual machines already configured (unless `--no-setup` option is given).

Growing a cluster (adding nodes to the cluster) should be supported by all the playbooks included in the elasticsearch package.

Basic usage of the command is:

```
usage: elasticsearch resize [-h] [-a N1:GROUP1[,N2:GROUP2]]
                             [-r N1:GROUP1[,N2:GROUP2]] [-v] [--no-setup]
                             [--yes]
                             cluster
```

`cluster` is the name of a cluster that has been *started* previously.

The following options are available:

**-h, --help** Show an help message and exits.

**-v, --verbose** Adding one or more `-v` will increase the verbosity accordingly.

**-a N1:GROUP1[,N2:GROUP2], --add N1:GROUP1[,N2:GROUP2]**

This option allow you to specify how many nodes for a specific group you want to add. You can specify multiple nodes separated by a comma.

Assuming you started, for instance, a cluster named *hadoop* using the default values stored in the configuration file:

```
hadoop-data_nodes=4
hadoop-task_nodes=4
```

and assuming you want to *add 5* more data nodes and *10* more task nodes, you can run:

```
elasticcluster resize -a 5:hadoop-data,10:hadoop-task
```

```
-r N1:GROUP1[,N2:GROUP2], --remove N1:GROUP1[,N2:GROUP2]
```

This option allow you to specify how many nodes you want to remove from a specific group. It follows the same syntax as the `--add` option.

**WARNING:** elasticcluster pick the nodes to remove at random, so **you have to be sure that any of the nodes can be removed**. Moreover, not all the playbooks support shrkinging!

```
--no-setup
```

By default elasticcluster will automatically run the **setup** command after starting and/or stopping the virtual machines. This option prevent the *setup* step to be run. **WARNING:** use this option wisely: depending on the cluster configuration it is impossible to know in advance what the status of the cluster will be after resizing it and NOT running the *setup* step.

```
--yes
```

Since resizing a cluster, especially shrinking, is a possibly disruptive action and is not supported by all the distributed playbooks, elasticcluster will always ask for confirmation before doing any modification, unless this option is given.

### 4.3.8 The `ssh` command

After a cluster is started, the easiest way to login on it is by using the **ssh** command. This command will run the *ssh* command with the correct options to connect to the cluster using the configured values for user and ssh key to use.

If no `ssh_to` option is specified in the configuration file, the **ssh** command will connect to the first host belonging to the type which comes first in alphabetic order, otherwise it will connect to the first host of the group specified by the `ssh_to` option of the `cluster` section. However, running the command `elasticcluster list-nodes <cluster>` will show which host will be used as frontend node.

The usage of the *ssh* command is as follow:

```
elasticcluster ssh <clustername> [ -- ssh arguments]
```

All the options and arguments following the `--` characters will be passed directly to the *ssh* command.

For instance, if you just want to run the `hostname -f` command on the frontend of the cluster you can run:

```
elasticcluster ssh <clustername> -- hostname -f
```

Note that since the IP address of the virtual machines are likely to be reused by different virtual machines, in order to avoid annoying warning messages from *ssh* elasticcluster will add the following options to the *ssh* command line:

- o **UserKnownHostsFile=/dev/null** Use an empty virtual file to check the host key of the remote machine.
- o **StrictHostKeyChecking=no** Disable check of the host key of the remove machine, without prompting to ask if the key can be accepted or not.

### 4.3.9 The `sftp` command

After a cluster is started, the easiest way to upload or download files to and from the cluster is by using the **sftp** command. This command will run the *sftp* command with the correct options to connect to the cluster using the configured values for user and ssh key to use.

If no `ssh_to` option is specified in the configuration file, the **sftp** command will connect to the first host belonging to the type which comes first in alphabetic order, otherwise it will connect to the first host of the group specified by

the `ssh_to` option of the `cluster` section. However, running the command `elasticcluster list-nodes <cluster>` will show which host will be used as frontend node.

The usage of the `sftp` command is as follow:

```
elasticcluster sftp <clustername> [ -- sftp arguments]
```

All the options and arguments following the `--` characters will be passed directly to the `sftp` command.

Note that since the IP address of the virtual machines are likely to be reused by different virtual machines, in order to avoid annoying warning messages from `ssh` elasticcluster will add the following options to the `sftp` command line:

- `-o UserKnownHostsFile=/dev/null` Use an empty virtual file to check the host key of the remote machine.
- `-o StrictHostKeyChecking=no` Disable check of the host key of the remove machine, without prompting to ask if the key can be accepted or not.

## 4.4 Playbooks distributed with elasticcluster

After the requested number of Virtual Machines have been started, elasticcluster uses [Ansible](#) to configure them based on the configuration options defined in the configuration file.

We distribute a few playbooks together with elasticcluster to configure some of the most wanted clusters. The playbooks are available at the `share/elasticcluster/providers/ansible-playbooks/` directory inside your virtualenv if you installed using `pip`, or in the `elasticcluster/providers/ansible-playbooks` directory of the github source code. You can copy, customize and redistribute them freely under the terms of the GPLv3 license.

A list of the most used playbooks distributed with elasticcluster and some explanation on how to use them follows.

### 4.4.1 Slurm

Tested on:

- Ubuntu 12.04
- Ubuntu 13.04
- Debian 7.1 (GCE)

ansible groups	role
<code>slurm_master</code>	Act as scheduler and submission host
<code>slurm_clients</code>	Act as compute node

This playbook will install the [SLURM](#) queue manager using the packages distributed with Ubuntu and will create a basic, working configuration.

You are supposed to only define one `slurm_master` and multiple `slurm_clients`. The first will act as login node and will run the scheduler, while the others will only execute the jobs.

The `/home` filesystem is exported *from* the slurm server to the compute nodes.

A *snippet* of a typical configuration for a slurm cluster is:

```
[cluster/slurm]
frontend_nodes=1
compute_nodes=5
ssh_to=frontend
setup_provider=ansible_slurm
...
```

```
[setup/ansible_slurm]
frontend_groups=slurm_master
compute_groups=slurm_clients
...
```

You can combine the slurm playbooks with ganglia. In this case the `setup` stanza will look like:

```
[setup/ansible_slurm]
frontend_groups=slurm_master,ganglia_master
compute_groups=slurm_clients,ganglia_monitor
...
```

## 4.4.2 Gridengine

Tested on:

- Ubuntu 12.04
- CentOS 6.3 (except for GCE images)
- Debian 7.1 (GCE)

ansible groups	role
gridengine_master	Act as scheduler and submission host
gridengine_clients	Act as compute node

This playbook will install [Grid Engine](#) using the packages distributed with Ubuntu or CentOS and will create a basic, working configuration.

You are supposed to only define one `gridengine_master` and multiple `gridengine_clients`. The first will act as login node and will run the scheduler, while the others will only execute the jobs.

The `/home` filesystem is exported *from* the gridengine server to the compute nodes. If you are running on a CentOS, also the `/usr/share/gridengine/default/common` directory is shared from the gridengine server to the compute nodes.

A *snippet* of a typical configuration for a gridengine cluster is:

```
[cluster/gridengine]
frontend_nodes=1
compute_nodes=5
ssh_to=frontend
setup_provider=ansible_gridengine
...

[setup/ansible_gridengine]
frontend_groups=gridengine_master
compute_groups=gridengine_clients
...
```

You can combine the gridengine playbooks with ganglia. In this case the `setup` stanza will look like:

```
[setup/ansible_gridengine]
frontend_groups=gridengine_master,ganglia_master
compute_groups=gridengine_clients,ganglia_monitor
...
```

Please note that Google Compute Engine provides Centos 6.2 images with a non-standard kernel which is **unsupported** by the gridengine packages.

### 4.4.3 HTCondor

Tested on:

- Ubuntu 12.04

ansible groups	role
condor_master	Act as scheduler, submission and execution host.
condor_workers	Act as execution host only.

This playbook will install the **HTCondor** workload management system using the packages provided by the Center for High Throughput Computing at UW-Madison.

The `/home` filesystem is exported *from* the condor master to the compute nodes.

A *snippet* of a typical configuration for a slurm cluster is:

```
[cluster/condor]
setup_provider=ansible_condor
frontend_nodes=1
compute_nodes=2
ssh_to=frontend
...

[setup/ansible_condor]
frontend_groups=condor_master
compute_groups=condor_workers
...
```

### 4.4.4 Ganglia

Tested on:

- Ubuntu 12.04
- CentOS 6.3
- Debian 7.1 (GCE)
- CentOS 6.2 (GCE)

ansible groups	role
ganglia_master	Run gmetad and web interface. It also run the monitor daemon.
ganglia_monitor	Run ganglia monitor daemon.

This playbook will install **Ganglia** monitoring tool using the packages distributed with Ubuntu or CentOS and will configure frontend and monitors.

You should run only one `ganglia_master`. This will install the `gmetad` daemon to collect all the metrics from the monitored nodes and will also run `apache`.

If the machine in which you installed `ganglia_master` has IP `10.2.3.4`, the ganglia web interface will be available at the address <http://10.2.3.4/ganglia/>

This playbook is supposed to be compatible with all the other available playbooks.

### 4.4.5 IPython cluster

Tested on:

- Ubuntu 12.04

- CentOS 6.3
- Debian 7.1 (GCE)
- CentOS 6.2 (GCE)

ansible groups	role
ipython_controller	Run an IPython cluster controller
ipython_engine	Run a number of ipython engine for each core

This playbook will install an **IPython cluster** to run python code in parallel on multiple machines.

One of the nodes should act as *controller* of the cluster (`ipython_controller`), running the both the *hub* and the *scheduler*. Other nodes will act as *engine*, and will run one “ipython engine” per core. You can use the *controller* node for computation too by assigning the `ipython_engine` class to it as well.

A *snippet* of typical configuration for an Hadoop cluster is:

```
[cluster/ipython]
setup_provider=ansible_ipython
controller_nodes=1
worker_nodes=4
ssh_to=controller
...

[setup/ansible_ipython]
controller_groups=ipython_controller,ipython_engine
worker_groups=ipython_engine
...
```

In order to use the IPython cluster, using the default configuration, you are supposed to connect to the controller node via ssh and run your code from there.

## 4.4.6 Hadoop

Tested on:

- Ubuntu 12.04
- CentOS 6.3
- Debian 7.1 (GCE)

ansible groups	role
hadoop_namenode	Run the Hadoop NameNode service
hadoop_jobtracker	Run the Hadoop JobTracker service
hadoop_datanode	Act as datanode for HDFS
hadoop_tasktracker	Act as tasktracker node accepting jobs from the JobTracker

Hadoop playbook will install a basic hadoop cluster using the packages available on the Hadoop website. The only supported version so far is **1.1.2 x86\_64** and it works both on CentOS and Ubuntu.

You must define only one `hadoop_namenode` and one `hadoop_jobtracker`. Configuration in which both roles belong to the same machines are not tested. You can mix `hadoop_datanode` and `hadoop_tasktracker` without problems though.

A *snippet* of a typical configuration for an Hadoop cluster is:

```
[cluster/hadoop]
hadoop-name_nodes=1
hadoop-jobtracker_nodes=1
```



```

hadoop-task-data_nodes=10
setup_provider=ansible_hadoop
ssh_to=hadoop-name
...

[setup/ansible_hadoop]
hadoop-name_groups=hadoop_namenode
hadoop-jobtracker_groups=hadoop_jobtracker
hadoop-task-data_groups=hadoop_tasktracker,hadoop_datanode
...
```

## 4.4.7 GlusterFS

Tested on:

- Ubuntu 12.04
- CentOS 6.3

ansible groups	role
gluster_data	Run a gluster <i>brick</i>
gluster_client	Install gluster client and install a gluster filesystem on /glusterfs

This will install a GlusterFS using all the `gluster_data` nodes as *bricks*, and any `gluster_client` to mount this filesystem in /glusterfs.

Setup is very basic, and by default no replicas is set.

To manage the gluster filesystem you need to connect to a `gluster_data` node.

## 4.4.8 OrangeFS/PVFS2

Tested on:

- Ubuntu 12.04

ansible groups	role
pvfs2_meta	Run the pvfs2 metadata service
pvfs2_data	Run the pvfs2 data node
pvfs2_client	configure as pvfs2 client and mount the filesystem

The OrangeFS/PVFS2 playbook will configure a pvfs2 cluster. It downloads the software from the [OrangeFS](#) website, compile and install it on all the machine, and run the various server and client daemons.

In addition, it will mount the filesystem in /pvfs2 on all the clients.

You can combine, for instance, a SLURM cluster with a PVFS2 cluster:

```

[cluster/slurm+pvfs2]
frontend_nodes=1
compute_nodes=10
pvfs2-nodes=10
ssh_to=frontend
setup_provider=ansible_slurm+pvfs2
...

[setup/ansible_slurm+pvfs2]
frontend_groups=slurm_master,pvfs2_client
compute_groups=slurm_clients,pvfs2_client
```

```
pvfs-nodes_groups=pvfs2_meta,pvfs2_data
...
```

This configuration will create a SLURM cluster with 10 compute nodes, 10 data nodes and a frontend, and will mount the `/pvfs2` directory from the data nodes to both the compute nodes and the frontend.

## 4.5 Elasicluster programming API

### 4.5.1 *elasticcluster*

#### Overview

Elasticcluster offers an API to programmatically manage compute clusters on cloud infrastructure. This page introduces the basic concepts of the API and provides sample code to illustrate the usage of the API. While this document should provide you with the basics, more details can be found in the respective module documentation

#### Getting Started

The following subchapters introduce the basic concepts of Elasticcluster.

#### Cluster

This is the heart of elasticcluster and handles all cluster relevant behavior. You can basically start, setup and stop a cluster. Also it provides factory methods to add nodes to the cluster. A typical workflow is as follows (see [slurm](#) for a code example):

1. create a new cluster
2. add nodes to fit your computing needs
3. start cluster; start all instances in the cloud
4. setup cluster; configure all nodes to fit your computing cluster
5. ssh into a node to submit computing jobs
6. eventually stop cluster; destroys all instances in the cloud

See documentation of the `Cluster` class for futher details.

#### Node

The node represents an instance in a cluster. It holds all information to connect to the nodes also manages the cloud instance. It provides the basic functionality to interact with the cloud instance, such as start, stop, check if the instance is up and ssh connect.

See the [Node](#) api docs for further details.

#### Cloud Provider

Manages the connection to the cloud webservice and offers all functionality used by the cluster to provision instances. Elasticcluster offers two different cloud providers at the current state:

- **BotoCloudProvider** Cloud provider to connect to EC2 compliant web services (e.g Amazon, Openstack, etc.)
- **GoogleCloudProvider** Cloud provider to connect to the Google Compute Engine (GCE)

All listed cloud providers above can be used to manage a cluster in the cloud. If the cloud operator is not supported by the implementations above, an alternative implementation can be provided by following the *AbstractCloudProvider* contract.

## Setup Provider

The setup provider configures in respect to the specified cluster and node configuration. The basic implementation *AnsibleSetupProvider* uses *ansible* to configure the nodes. Ansible is a push based configuration management system in which the configuration is stored locally and pushed to all the nodes in the cluster.

See the *Playbooks distributed with elasticsearch* page for more details on the cluster setups possible with the ansible implementation and how the ansible playbooks can be enhanced.

If this implementation does not satisfy the clients needs, an alternative implementation can be implemented following the *AbstractSetupProvider* contract.

## Cluster Repository

The cluster repository is responsible to keep track of multiple clusters over time. Therefore Elasticsearch provides two implementations:

- **MemRepository** Stores the clusters in memory. Therefore after stopping a program using this repository, all clusters are not recoverable but possibly still running.
- **ClusterRepository** Stores the cluster on disk persistently. This implementation uses pickle to serialize and deserialize the cluster.

If a client wants to store the cluster in a database for example, an alternative implementation can be provided following the *AbstractClusterRepository* contract.

## Sample Code

### Start and setup a SLURM cluster

The following sample code shows how to start and setup a SLURM cluster on an OpenStack cloud and provides further information on each step. Other cluster types on other cloud providers can be setup accordingly.

```
import elasticsearch

# Initialise an EC2 compatible cloud provider, in this case an OpenStack
# cloud operator is chosen. To initialise the cloud provider the
# following parameters are passed:
#   url:          url to connect to the cloud operator web service
#   region:       region to start the nodes on
#   access_key:   access key of the current user to connect
#   secret_key:   secret key of the current user to connect
cloud_provider = elasticsearch.BotoCloudProvider(
    'http://uzh.ch/services/Cloud',
    'nova', 'access_key', 'secret_key')

# Initialising the setup provider needs a little more preparation:
```

```
# the groups dictionary specifies the kind of nodes used for this cluster.
# In this case we want a frontend and a compute kind. The frontend node
# (s) will be setup as slurm_master, the compute node(s) as slurm_clients.
# This corresponds to the documentation of the ansible playbooks
# provided with elasticcluster. The kind of the node is a name specified
# by the user. This name will be used to set a new hostname on the
# instance, therefore it should meet the requirements of RFC 953
# groups['kind'] = ['ansible_group1', 'ansible_group2']
groups = dict()
groups['frontend'] = ['slurm_master']
groups['compute'] = ['slurm_clients']

setup_provider = elasticcluster.AnsibleSetupProvider(groups)

# cluster initialisation (note: ssh keys are same for all nodes)
# After the steps above initialising an empty cluster is a piece of cake.
# The cluster takes the following arguments:
# name:          name to identify the cluster
# cloud_provider: cloud provider to connect to cloud
# setup_provider: setup provider to configure the cluster
# ssh_key_name:  name of the ssh key stored (or to be stored) on the
#               cloud
# ssh_key_pub:   path to public ssh key file
# ssh_key_priv:  path to private ssh key file
#
# The ssh key files are used for all instances in this cluster.
cluster = elasticcluster.Cluster('my-cluster', cloud_provider,
                                setup_provider, 'ssh_key_name',
                                '~/ssh/keys/my_ssh_key.pub',
                                '~/ssh/keys/my_ssh_key')

# To add nodes to the cluster we can use the add_node. This
# only initialises a new node, but does not start it yet.
# The add node function is basically a factory method to make it easy to
# add nodes to a cluster. It takes the following arguments:
# kind:  kind of the node in this cluster. This corresponds to the
#        groups defined in the cloud_provider.
cluster.add_node('frontend', 'ami-00000048', 'gc3-user',
                'm1.tiny', 'all_tcp_ports')

# We can also add multiple nodes with the add_nodes method.
# The following command will add 2 nodes of the kind `compute` to the
# cluster
cluster.add_nodes('compute', 2, 'ami-00000048', 'gc3-user', 'm1.tiny',
                'all_tcp_ports')

# Since we initialised all the nodes for this computing cluster,
# we can finally start the cluster.
# The start method is blocking and does the following tasks:
# * call the cloud provider to start an instance for each node in a
#   separate thread.
# * to make sure elasticcluster is not stopped during creation of an
#   instance, it will overwrite the sigint handler
# * waits until all nodes are alive (meaning ssh connection
#   works)
# * If the startup timeout is reached and not all nodes are alive,
#   the cluster will stop and destroy all instances
cluster.start()
```

```
# Now, all the nodes are started and we can call the setup method to
# configure slurm on the nodes.
cluster.setup()
```

### Asynchronous node start

The `start()` method of the `Cluster()` class is blocking and therefore waits until all nodes are alive. If a client wants to use this time for other tasks, the nodes can as well be started asynchronous:

```
# retrieve all nodes from the cluster
nodes = cluster.get_all_nodes()

# start each node
# The start method on the node is non blocking and will return as soon
# as the cloud provider is contacted to start a new instance
for node in nodes:
    node.start()

# wait until all nodes are alive
starting_nodes = nodes[:]
while starting_nodes:
    starting_nodes = [n for n in starting_nodes if not n.is_alive()]
```

### Storing a cluster on disk

By default elasticcluster will store the cluster in memory only. Therefore after a programm shutdown the cluster will not be available anymore in elasticcluster, but might still be running on the cloud. The following example shows how to store clusters on disk to retrieve after a programm restart:

```
# The cluster repository uses pickle to store clusters each in a
# separate file in the provided storage directory.
repository = elasticcluster.ClusterRepository('/path/to/storage/dir')

# On cluster initialisation we can pass the repository as optional
# argument.
cluster = elasticcluster.Cluster('my-cluster', cloud_provider,
                                setup_provider, 'ssh_key_name',
                                '~/ssh/keys/my_ssh_key.pub',
                                '~/ssh/keys/my_ssh_key',
                                repository=repository)

# When starting the cluster, it will save its state using the repository.
cluster.start()
```

After a program shutdown we can therefore fetch the cluster from the repository again and work with it as expected:

```
repository = elasticcluster.ClusterRepository('/path/to/storage/dir')

# retrieve the cluster from the repository
cluster = repository.get('my-cluster')

# or retrieve all clusters that are stored in the repository
clusters = repository.get_all()
```

## Logging

Elasticsearch uses the python *logging* module to log events. A client can overwrite the settings as illustrated below:

```
import logging

import elasticsearch

log = elasticsearch.log
level = logging.getLevelName('INFO')
log.setLevel(level)
```

The current example only shows how to increase the log level, but any settings can be applied compliant with the logging module of python.

### 4.5.2 *elasticsearch.cluster*

**class** `elasticsearch.cluster.Cluster`(*name, cloud\_provider, setup\_provider, user\_key\_name, user\_key\_public, user\_key\_private, repository=None, \*\*extra*)

This is the heart of elasticsearch and handles all cluster relevant behavior. You can basically start, setup and stop a cluster. Also it provides factory methods to add nodes to the cluster. A typical workflow is as follows:

- create a new cluster
- add nodes to fit your computing needs
- start cluster; start all instances in the cloud
- setup cluster; configure all nodes to fit your computing cluster
- eventually stop cluster; destroys all instances in the cloud

#### Parameters

- **name** (*str*) – unique identifier of the cluster
- **cloud\_provider** (`elasticsearch.providers.AbstractCloudProvider`) – access to the cloud to manage nodes
- **setup\_provider** (`elasticsearch.providers.AbstractSetupProvider`) – provider to setup cluster
- **user\_key\_name** (*str*) – name of the ssh key to connect to cloud
- **user\_key\_public** (*str*) – path to ssh public key file
- **user\_key\_private** (*str*) – path to ssh private key file
- **repository** (`elasticsearch.repository.AbstractClusterRepository`) – by default the `elasticsearch.repository.MemoryRepository` is used to store the cluster in memory. Provide another repository to store the cluster in a persistent state.
- **extra** – tbd.

**Variables** **nodes** – dict [*node\_type*] = [*Node*] that represents all nodes in this cluster

**add\_node** (*kind, image\_id, image\_user, flavor, security\_group, image\_userdata='', name=None, \*\*extra*)

Adds a new node to the cluster. This factory method provides an easy way to add a new node to the cluster

by specifying all relevant parameters. The node does not get started nor setup automatically, this has to be done manually afterwards.

#### Parameters

- **kind** (*str*) – kind of node to start. this refers to the groups defined in the ansible setup provider `elasticcluster.providers.AnsibleSetupProvider` Please note that this must match the `[a-zA-Z0-9-]` regexp, as it is used to build a valid hostname
- **image\_id** (*str*) – image id to use for the cloud instance (e.g. ami on amazon)
- **image\_user** (*str*) – user to login on given image
- **flavor** (*str*) – machine type to use for cloud instance
- **security\_group** (*str*) – security group that defines firewall rules to the instance
- **image\_userdata** (*str*) – commands to execute after instance starts
- **name** (*str*) – name of this node, automatically generated if None

**Raises** `ValueError`: *kind* argument is an invalid string.

**Returns** created `Node`

**add\_nodes** (*kind, num, image\_id, image\_user, flavor, security\_group, image\_userdata=''*, *\*\*extra*)  
 Helper method to add multiple nodes of the same kind to a cluster.

#### Parameters

- **kind** (*str*) – kind of node to start. this refers to the groups defined in the ansible setup provider `elasticcluster.providers.AnsibleSetupProvider`
- **num** (*int*) – number of nodes to add of this kind
- **image\_id** (*str*) – image id to use for the cloud instance (e.g. ami on amazon)
- **image\_user** (*str*) – user to login on given image
- **flavor** (*str*) – machine type to use for cloud instance
- **security\_group** (*str*) – security group that defines firewall rules to the instance
- **image\_userdata** (*str*) – commands to execute after instance starts

**get\_all\_nodes** ()

Returns a list of all nodes in this cluster as a mixed list of different node kinds.

**Returns** list of `Node`

**get\_frontend\_node** ()

Returns the first node of the class specified in the configuration file as *ssh\_to*, or the first node of the first class in alphabetic order.

**Returns** `Node`

**Raise** `elasticcluster.exceptions.NodeNotFound` if no valid frontend node is found

**remove\_node** (*node*)

Removes a node from the cluster, but does not stop it. Use this method with caution.

**Parameters** **node** (`Node`) – node to remove

**setup** ()

Configure the cluster nodes with the specified This is delegated to the provided `elasticcluster.providers.AbstractSetupProvider`

**Returns** bool - True on success, False otherwise

**start** (*min\_nodes=None*)

Starts up all the instances in the cloud. To speed things up all instances are started in a separate thread. To make sure elasticcluster is not stopped during creation of an instance, it will overwrite the sigint handler. As soon as the last started instance is returned and saved to the repository, sigint is executed as usual. An instance is up and running as soon as a ssh connection can be established. If the startup timeout is reached before all instances are started, the cluster will stop and destroy all instances.

This method is blocking and might take some time depending on the amount of instances to start.

**Parameters** **min\_nodes** (*dict [node\_kind] = number*) – minimum number of nodes to start in case the quota is reached before all instances are up

**startup\_timeout = 600**

timeout in seconds to start all nodes

**stop** (*force=False*)

Destroys all instances of this cluster and calls delete on the repository.

**Parameters** **force** (*bool*) – force termination of instances in any case

**update** ()

Update all connection information of the nodes of this cluster. It occurs for example public ip's are not available immediately, therefore calling this method might help.

**class** elasticcluster.cluster.**Node** (*name, cluster\_name, kind, cloud\_provider, user\_key\_public, user\_key\_private, user\_key\_name, image\_user, security\_group, image, flavor, image\_userdata=None, \*\*extra*)

The node represents an instance in a cluster. It holds all information to connect to the nodes also manages the cloud instance. It provides the basic functionality to interact with the cloud instance, such as start, stop, check if the instance is up and ssh connect.

#### Parameters

- **name** (*str*) – identifier of the node
- **kind** (*str*) – kind of node in regard to cluster. this usually refers to a specified group in the *elasticcluster.providers.AbstractSetupProvider*
- **cloud\_provider** (*elasticcluster.providers.AbstractCloudProvider*) – cloud provider to manage the instance
- **user\_key\_public** (*str*) – path to the ssh public key
- **user\_key\_private** (*str*) – path to the ssh private key
- **user\_key\_name** (*str*) – name of the ssh key
- **image\_user** (*str*) – user to connect to the instance via ssh
- **security\_group** (*str*) – security group to setup firewall rules
- **image** (*str*) – image id to launch instance with
- **flavor** (*str*) – machine type to launch instance
- **image\_userdata** (*str*) – commands to execute after instance start

#### Variables

- **instance\_id** – id of the node instance on the cloud
- **preferred\_ip** – IP address used to connect to the node.
- **ips** – list of all the IPs defined for this node.



**connect ()**

Connect to the node via ssh using the paramiko library.

**Returns** `paramiko.SSHClient` - ssh connection or `None` on failure

**connection\_ip ()**

Returns the IP to be used to connect to this node.

If the instance has a public IP address, then this is returned, otherwise, its private IP is returned.

**connection\_timeout = 5**

timeout in seconds to connect to host via ssh

**is\_alive ()**

Checks if the current node is up and running in the cloud. It only checks the status provided by the cloud interface. Therefore a node might be running, but not yet ready to ssh into it.

**pprint ()**

Pretty print information about the node.

**Returns** `str` - representaion of a node in pretty print

**start ()**

Starts the node on the cloud using the given instance properties. This method is non-blocking, as soon as the node id is returned from the cloud provider, it will return. Therefore the `is_alive` and `update_ips` methods can be used to further gather details about the state of the node.

**stop ()**

Destroys the instance launched on the cloud for this specific node.

**update\_ips ()**

Retrieves the public and private ip of the instance by using the cloud provider. In some cases the public ip assignment takes some time, but this method is non blocking. To check for a public ip, consider calling this method multiple times during a certain timeout.

### 4.5.3 *elasticcluster.conf*

**class** `elasticcluster.conf.ConfigReader (configfile)`

Reads the configuration properties from a ini file.

**Parameters** `configfile (str)` – path to configfile

**read\_config ()**

Reads the configuration properties from the ini file and links the section to comply with the cluster config dictionary format.

**Returns** dictionary containing all configuration properties from the ini file in compliance to the cluster config format

**Raises** `voluptuous.MultipleInvalid` if not all sections present or broken links between scitons

**class** `elasticcluster.conf.ConfigValidator (config)`

Validator for the cluster configuration dictionary.

**Parameters** `config` – dictionary containing cluster configuration properties

**validate ()**

Validates the given configuration `self.config` to comply with elasticcluster. As well all types are converted to the expected format if possible.

**Raises** `voluptuous.MultipleInvalid` if multiple properties are not compliant

**Raises** `voluptuous.Invalid` if one property is invalid

**class** `elasticcluster.conf.Configurator` (*cluster\_conf*, *storage\_path=None*)

The **Configurator** is responsible for (I) keeping track of the configuration and (II) offer factory methods to create all kind of objects that need information from the configuration.

The cluster configuration dictionary is structured in the following way: (see an example @

<https://github.com/gc3-uzh-ch/elasticcluster/wiki/Configuration-Module>)

```
{ "<cluster_template>" : {
    "setup" : { properties of the setup section },
    "cloud" : { properties of the cloud section },
    "login" : { properties of the login section },
    "cluster" : { properties of the cluster section },
    "nodes": {  "<node_kind>" : { properties of the node},
                "<node_kind>" : { properties of the node},
            },
    },
    "<cluster_template>" : {
        (see above)
    }
}
```

#### Parameters

- **cluster\_conf** (*dict*) – see description above
- **storage\_path** (*str*) – path to store data

**Raises** `MultipleInvalid` – configuration validation

**create\_cloud\_provider** (*cluster\_template*)

Creates a cloud provider by inspecting the configuration properties of the given cluster template.

**Parameters** **cluster\_template** (*str*) – template to use (if not already specified on init)

**Returns** cloud provider that fulfills the contract of `elasticcluster.providers.AbstractSetupProvider`

**create\_cluster** (*template*, *name=None*)

Creates a cluster by inspecting the configuration properties of the given cluster template.

#### Parameters

- **template** (*str*) – name of the cluster template
- **name** (*str*) – name of the cluster. If not defined, the cluster

will be named after the template.

**Returns** `elasticcluster.cluster.cluster` instance

**Raises** `ConfigurationError` – cluster template not found in config

**create\_setup\_provider** (*cluster\_template*, *name=None*)

Creates the setup provider for the given cluster template.

#### Parameters

- **cluster\_template** (*str*) – template of the cluster
- **name** (*str*) – name of the cluster to read configuration properties

**classmethod** **fromConfig** (*configfile*, *storage\_path=None*)

Helper method to initialize Configurator from an ini file.

**Parameters** `configfile` (*str*) – path to the ini file

**Returns** *Configurator*

**load\_cluster** (*cluster\_name*)

Loads a cluster from the cluster repository.

**Parameters** `cluster_name` (*str*) – name of the cluster

**Returns** `elasticcluster.cluster.cluster` instance

#### 4.5.4 *elasticcluster.exceptions*

#### 4.5.5 *elasticcluster.gc3\_config*

#### 4.5.6 *elasticcluster.helpers*

**class** `elasticcluster.helpers.Singleton` (*decorated*)

see: <http://stackoverflow.com/a/7346105> A non-thread-safe helper class to ease implementing singletons. This should be used as a decorator – not a metaclass – to the class that should be a singleton.

The decorated class can define one `__init__` function that takes only the *self* argument. Other than that, there are no restrictions that apply to the decorated class.

To get the singleton instance, use the *Instance* method. Trying to use `__call__` will result in a *TypeError* being raised.

Limitations: The decorated class cannot be inherited from.

**Instance** ()

Returns the singleton instance. Upon its first call, it creates a new instance of the decorated class and calls its `__init__` method. On all subsequent calls, the already created instance is returned.

#### 4.5.7 *elasticcluster.main*

#### 4.5.8 *elasticcluster.module*

#### 4.5.9 *elasticcluster.providers*

**class** `elasticcluster.providers.AbstractCloudProvider` (*\*\*config*)

Defines the contract for a cloud provider to proper function with elasticcluster.

**get\_ips** (*instance\_id*)

Retrieves the private and public ip addresses for a given instance.

**Returns** list (IPs)

**is\_instance\_running** (*instance\_id*)

Checks if the instance is up and running.

**Parameters** `instance_id` (*str*) – instance identifier

**Returns** bool - True if running, False otherwise

**start\_instance** (*key\_name*, *public\_key\_path*, *private\_key\_path*, *security\_group*, *flavor*, *image\_id*, *image\_userdata*, *username=None*, *node\_name=None*)

Starts a new instance on the cloud using the given properties. Multiple instances might be started in different threads at the same time. The implementation should handle any problems regarding this itself.

#### Parameters

- **key\_name** (*str*) – name of the ssh key to connect
- **public\_key\_path** (*str*) – path to ssh public key
- **private\_key\_path** (*str*) – path to ssh private key
- **security\_group** (*str*) – firewall rule definition to apply on the instance
- **flavor** (*str*) – machine type to use for the instance
- **image\_name** (*str*) – image type (os) to use for the instance
- **image\_userdata** (*str*) – command to execute after startup
- **username** (*str*) – username for the given ssh key, default None

**Returns** *str* - instance id of the started instance

**stop\_instance** (*instance\_id*)

Stops the instance gracefully.

**Parameters** **instance\_id** (*str*) – instance identifier

**Returns** None

**class** elasticcluster.providers.**AbstractSetupProvider**

TODO: define...

**cleanup** ()

Cleanup any temporary file or directory created during setup. This method is called every time a cluster is stopped.

**Returns** None

**setup\_cluster** (*cluster*)

Configures all nodes of a cluster to function in respect to the given configuration.

This method *must* be idempotent, i.e. it should always be safe calling it multiple times..

**Parameters** **cluster** (*elasticcluster.cluster.Cluster*) – cluster to configure

**Returns** *True* if the cluster is correctly configured, even if the method didn't actually do anything. *False* if the cluster is not configured.

### 4.5.10 elasticcluster.providers.ansible\_provider

```
class elasticcluster.providers.ansible_provider.AnsibleSetupProvider (groups,
                                                                    play-
                                                                    book_path=None,
                                                                    environ-
                                                                    ment_vars={},
                                                                    stor-
                                                                    age_path=None,
                                                                    sudo=True,
                                                                    sudo_user='root',
                                                                    ansi-
                                                                    ble_module_dir=None,
                                                                    **ex-
                                                                    tra_conf)
```

This implementation uses ansible to configure and manage the cluster setup. See <https://github.com/ansible/ansible> for details.

### Parameters

- **groups** (*dict*) – dictionary of node kinds with corresponding ansible groups to install on the node kind. e.g [node\_kind] = ['ansible\_group1', 'ansible\_group2'] The group defined here can be references in each node. Therefore groups can make it easier to define multiple groups for one node.
- **playbook\_path** (*str*) – path to playbook; if empty this will use the shared playbook of elasticcluster
- **environment\_vars** (*dict*) – dictionary to define variables per node kind, e.g. [node\_kind][var] = value
- **storage\_path** (*str*) – path to store the inventory file. By default the inventory file is saved temporarily in a temporary directory and deleted when the cluster is stopped.
- **sudo** (*bool*) – indication whether use sudo to gain root permission
- **sudo\_user** (*str*) – user with root permission
- **ansible\_module\_dir** (*str*) – path to addition ansible modules
- **extra\_conf** – tbd.

### Variables

- **groups** – node kind and ansible group mapping dictionary
- **environment** – additional environment variables

### **cleanup** (*cluster*)

Deletes the inventory file used last recently used.

**Parameters** **cluster** (*elasticcluster.cluster.Cluster*) – cluster to clear up inventory file for

### **setup\_cluster** (*cluster*)

Configures the cluster according to the node\_kind to ansible group matching. This method is idempotent and therefore can be called multiple times without corrupting the cluster configuration.

**Parameters** **cluster** (*elasticcluster.cluster.Cluster*) – cluster to configure

**Returns** True on success, False otherwise. Please note, if nothing has to be configured True is returned

**Raises** *AnsibleError* if the playbook can not be found or playbook is corrupt

## 4.5.11 *elasticcluster.providers.ec2\_boto*

```
class elasticcluster.providers.ec2_boto.BotoCloudProvider (ec2_url,          ec2_region,
                                                         ec2_access_key,
                                                         ec2_secret_key,          stor-
                                                         age_path=None,          re-
                                                         quest_floating_ip=False)
```

This implementation of *elasticcluster.providers.AbstractCloudProvider* uses the boto ec2 interface to connect to ec2 compliant clouds and manage instances.

Please check <https://github.com/boto/boto> for further information about the supported cloud platforms.

### Parameters

- **ec2\_url** (*str*) – url to connect to cloud web service
- **ec2\_region** (*str*) – region identifier

- **ec2\_access\_key** (*str*) – access key of the user account
- **ec2\_secret\_key** (*str*) – secret key of the user account
- **storage\_path** (*str*) – path to store temporary data
- **request\_floating\_ip** (*bool*) – Whether ip are assigned automatically *True* or floating ips have to be assigned manually *False*

**get\_ips** (*instance\_id*)

Retrieves the private and public ip addresses for a given instance.

**Returns** list (ips)

**is\_instance\_running** (*instance\_id*)

Checks if the instance is up and running.

**Parameters** **instance\_id** (*str*) – instance identifier

**Returns** bool - True if running, False otherwise

**start\_instance** (*key\_name*, *public\_key\_path*, *private\_key\_path*, *security\_group*, *flavor*, *image\_id*, *image\_userdata*, *username=None*, *node\_name=None*, *\*\*kwargs*)

Starts a new instance on the cloud using the given properties. The following tasks are done to start an instance:

- establish a connection to the cloud web service
- check ssh keypair and upload it if it does not yet exist. This is a locked process, since this function might be called in multiple threads and we only want the key to be stored once.
- check if the security group exists
- run the instance with the given properties

**Parameters**

- **key\_name** (*str*) – name of the ssh key to connect
- **public\_key\_path** (*str*) – path to ssh public key
- **private\_key\_path** (*str*) – path to ssh private key
- **security\_group** (*str*) – firewall rule definition to apply on the instance
- **flavor** (*str*) – machine type to use for the instance
- **image\_id** (*str*) – image type (os) to use for the instance
- **image\_userdata** (*str*) – command to execute after startup
- **username** (*str*) – username for the given ssh key, default None

**Returns** str - instance id of the started instance

**stop\_instance** (*instance\_id*)

Stops the instance gracefully.

**Parameters** **instance\_id** (*str*) – instance identifier

#### 4.5.12 *elasticcluster.providers.gce*

Cloud provider for the Google Compute Engine.

See <https://code.google.com/p/google-cloud-platform-samples/source/browse/python-client-library-example/gce.py?repo=compute> for reference.

`elasticcluster.providers.gce.GCE_SCOPE = 'https://www.googleapis.com/auth/compute'`  
the OAuth scope for the GCE web API

```
class elasticcluster.providers.gce.GoogleCloudProvider(gce_client_id, gce_client_secret,
                                                       gce_project_id,      zone='us-
central1-a',  network='default',
                                                       email='default',      stor-
age_path=None)
```

Cloud provider for the Google Compute Engine.

#### Parameters

- **gce\_client\_id** (*str*) – Client ID to use in OAuth authentication.
- **gce\_client\_secret** (*str*) – Client secret (password) to use in OAuth authentication.
- **gce\_project\_id** (*str*) – Project name to log in to GCE.
- **zone** (*str*) – gce zone, default is *us-central1-a*
- **network** (*str*) – network to use, default is *default*
- **email** (*str*) – service email to use, default is *default*
- **storage\_path** (*str*) – path to store authentication data (oauth.dat file). If no path is specified, the login data has to be entered after every request.

**get\_ips** (*instance\_id*)

Retrieves the ip addresses (private and public) from the cloud provider by the given instance id.

**Parameters** **instance\_id** (*str*) – id of the instance

**Returns** list (ips)

**Raises** InstanceError if the ip could not be retrieved.

**is\_instance\_running** (*instance\_id*)

Check whether the instance is up and running.

**Parameters** **instance\_id** (*str*) – instance identifier

**Reutrn** True if instance is running, False otherwise

**list\_instances** (*filter=None*)

List instances on GCE, optionally filtering the results.

**Parameters** **filter** (*str*) – Filter specification; see <https://developers.google.com/compute/docs/reference/latest/instances/list> for details.

**Returns** list of instances

**start\_instance** (*key\_name, public\_key\_path, private\_key\_path, security\_group, flavor, image\_id, image\_userdata, username=None, instance\_name=None, \*\*kwargs*)

Starts a new instance with the given properties and returns the instance id.

#### Parameters

- **key\_name** (*str*) – name of the ssh key to connect
- **public\_key\_path** (*str*) – path to ssh public key
- **private\_key\_path** (*str*) – path to ssh private key
- **security\_group** (*str*) – firewall rule definition to apply on the instance

- **flavor** (*str*) – machine type to use for the instance
- **image\_id** (*str*) – image type (os) to use for the instance
- **image\_userdata** (*str*) – command to execute after startup
- **username** (*str*) – username for the given ssh key, default None
- **instance\_name** (*str*) – name of the instance

**Returns** *str* - instance id of the started instance

**stop\_instance** (*instance\_id*)

Stops the instance gracefully.

**Parameters** **instance\_id** (*str*) – instance identifier

**Raises** *InstanceError* if instance can not be stopped

### 4.5.13 *elasticcluster.repository*

**class** *elasticcluster.repository.AbstractClusterRepository*

Defines the contract for a cluster repository to store clusters in a persistent state.

**delete** (*cluster*)

Deletes the cluster from persistent state.

**Parameters** **cluster** (*elasticcluster.cluster.Cluster*) – cluster to delete from persistent state

**get** (*name*)

Retrieves the cluster by the given name.

**Parameters** **name** (*str*) – name of the cluster (identifier)

**Returns** instance of *elasticcluster.cluster.Cluster* that matches the given name

**get\_all** ()

Retrieves all stored clusters from the persistent state.

**Returns** list of *elasticcluster.cluster.Cluster*

**save\_or\_update** (*cluster*)

Save or update the cluster in a persistent state. Elasticcluster will call this method multiple times, so the implementation should handle save and update seamlessly

**Parameters** **cluster** (*elasticcluster.cluster.Cluster*) – cluster object to store

**class** *elasticcluster.repository.ClusterRepository* (*storage\_path*)

This implementation of *AbstractClusterRepository* stores the cluster on the local disc using pickle. Therefore the cluster object and all its dependencies will be saved in a pickle (binary) file.

**Parameters** **storage\_path** (*str*) – path to the folder to store the cluster information

**delete** (*cluster*)

Deletes the cluster from persistent state.

**Parameters** **cluster** (*elasticcluster.cluster.Cluster*) – cluster to delete from persistent state

**get** (*name*)

Retrieves the cluster with the given name.

**Parameters** **name** (*str*) – name of the cluster (identifier)



**Returns** *elasticcluster.cluster.Cluster*

**get\_all()**

Retrieves all clusters from the persistent state.

**Returns** list of *elasticcluster.cluster.Cluster*

**save\_or\_update(cluster)**

Save or update the cluster to persistent state.

**Parameters cluster** (*elasticcluster.cluster.Cluster*) – cluster to save or update

**class** *elasticcluster.repository.MemRepository*

This implementation of *AbstractClusterRepository* stores the clusters in memory, without actually saving the data on disk.

**delete(cluster)**

Deletes the cluster from memory.

**Parameters cluster** (*elasticcluster.cluster.Cluster*) – cluster to delete

**get(name)**

Retrieves the cluster by the given name.

**Parameters name** (*str*) – name of the cluster (identifier)

**Returns** instance of *elasticcluster.cluster.Cluster* that matches the given name

**get\_all()**

Retrieves all stored clusters from the memory.

**Returns** list of *elasticcluster.cluster.Cluster*

**save\_or\_update(cluster)**

Save or update the cluster in a memory.

**Parameters cluster** (*elasticcluster.cluster.Cluster*) – cluster object to store

#### 4.5.14 *elasticcluster.subcommands*

**class** *elasticcluster.subcommands.AbstractCommand(params)*

Defines the general contract every command has to fulfill in order to be recognized by the arguments list and executed afterwards.

**execute()**

This method is executed after a command was recognized and may vary in its behavior.

**pre\_run()**

Overrides this method to execute any pre-run code, especially to check any command line options.

**setup(subparsers)**

This method handles the setup of the subcommand. In order to do so, every command has to add a parser to the subparsers reference given as parameter. The following example is the minimum implementation of such a setup procedure: `parser = subparsers.add_parser("start") parser.set_defaults(func=self.execute)`

**class** *elasticcluster.subcommands.ListClusters(params)*

Print a list of all clusters that have been started.

**class** *elasticcluster.subcommands.ListNodes(params)*

Show some information on all the nodes belonging to a given cluster.

**execute()**

Lists all nodes within the specified cluster with certain information like id and ip.

```
class elasticcluster.subcommands.ListTemplates (params)
    List the available templates defined in the configuration file.

class elasticcluster.subcommands.ResizeCluster (params)
    Resize the cluster by adding or removing compute nodes.

class elasticcluster.subcommands.SetupCluster (params)
    Setup the given cluster by calling the setup provider defined for this cluster.

class elasticcluster.subcommands.SftpFrontend (params)
    Open an SFTP session to the cluster frontend host.

class elasticcluster.subcommands.SshFrontend (params)
    Connect to the frontend of the cluster using ssh.

class elasticcluster.subcommands.Start (params)
    Create a new cluster using the given cluster template.

    execute ()
        Starts a new cluster.

class elasticcluster.subcommands.Stop (params)
    Stop a cluster and terminate all associated virtual machines.

    execute ()
        Stops the cluster if it's running.

    setup (subparsers)
        @see abstract_command contract
```

---

## Indices and tables

---

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



## e

- [elasticcluster](#), 34
- [elasticcluster.cluster](#), 34
- [elasticcluster.conf](#), 37
- [elasticcluster.exceptions](#), 39
- [elasticcluster.helpers](#), 39
- [elasticcluster.main](#), 39
- [elasticcluster.module](#), 39
- [elasticcluster.providers](#), 39
  - [elasticcluster.providers.ansible\\_provider](#), 40
- [elasticcluster.providers.ec2\\_boto](#), 41
- [elasticcluster.providers.gce](#), 42
- [elasticcluster.repository](#), 44
- [elasticcluster.subcommands](#), 45



## A

AbstractCloudProvider (class in elasticcluster.providers), 39  
AbstractClusterRepository (class in elasticcluster.repository), 44  
AbstractCommand (class in elasticcluster.subcommands), 45  
AbstractSetupProvider (class in elasticcluster.providers), 40  
add\_node() (elasticcluster.cluster.Cluster method), 34  
add\_nodes() (elasticcluster.cluster.Cluster method), 35  
AnsibleSetupProvider (class in elasticcluster.providers.ansible\_provider), 40

## B

BotoCloudProvider (class in elasticcluster.providers.ec2\_boto), 41

## C

cleanup() (elasticcluster.providers.AbstractSetupProvider method), 40  
cleanup() (elasticcluster.providers.ansible\_provider.AnsibleSetupProvider method), 41  
Cluster (class in elasticcluster.cluster), 34  
ClusterRepository (class in elasticcluster.repository), 44  
ConfigReader (class in elasticcluster.conf), 37  
Configurator (class in elasticcluster.conf), 38  
ConfigValidator (class in elasticcluster.conf), 37  
connect() (elasticcluster.cluster.Node method), 36  
connection\_ip() (elasticcluster.cluster.Node method), 37  
connection\_timeout (elasticcluster.cluster.Node attribute), 37  
create\_cloud\_provider() (elasticcluster.conf.Configurator method), 38  
create\_cluster() (elasticcluster.conf.Configurator method), 38  
create\_setup\_provider() (elasticcluster.conf.Configurator method), 38

## D

delete() (elasticcluster.repository.AbstractClusterRepository

method), 44

delete() (elasticcluster.repository.ClusterRepository method), 44

delete() (elasticcluster.repository.MemRepository method), 45

## E

elasticcluster (module), 34  
elasticcluster.cluster (module), 34  
elasticcluster.conf (module), 37  
elasticcluster.exceptions (module), 39  
elasticcluster.helpers (module), 39  
elasticcluster.main (module), 39  
elasticcluster.module (module), 39  
elasticcluster.providers (module), 39  
elasticcluster.providers.ansible\_provider (module), 40  
elasticcluster.providers.ec2\_boto (module), 41  
elasticcluster.providers.gce (module), 42  
elasticcluster.repository (module), 44  
elasticcluster.subcommands (module), 45  
execute() (elasticcluster.subcommands.AbstractCommand method), 45  
execute() (elasticcluster.subcommands.ListNodes method), 45  
execute() (elasticcluster.subcommands.Start method), 46  
execute() (elasticcluster.subcommands.Stop method), 46

## F

fromConfig() (elasticcluster.conf.Configurator class method), 38

## G

GCE\_SCOPE (in module elasticcluster.providers.gce), 43  
get() (elasticcluster.repository.AbstractClusterRepository method), 44  
get() (elasticcluster.repository.ClusterRepository method), 44  
get() (elasticcluster.repository.MemRepository method), 45  
get\_all() (elasticcluster.repository.AbstractClusterRepository method), 44

get\_all() (elasticcluster.repository.ClusterRepository method), 45  
 get\_all() (elasticcluster.repository.MemRepository method), 45  
 get\_all\_nodes() (elasticcluster.cluster.Cluster method), 35  
 get\_frontend\_node() (elasticcluster.cluster.Cluster method), 35  
 get\_ips() (elasticcluster.providers.AbstractCloudProvider method), 39  
 get\_ips() (elasticcluster.providers.ec2\_boto.BotoCloudProvider method), 42  
 get\_ips() (elasticcluster.providers.gce.GoogleCloudProvider method), 43  
 GoogleCloudProvider (class in elasticcluster.providers.gce), 43

## I

Instance() (elasticcluster.helpers.Singleton method), 39  
 is\_alive() (elasticcluster.cluster.Node method), 37  
 is\_instance\_running() (elasticcluster.providers.AbstractCloudProvider method), 39  
 is\_instance\_running() (elasticcluster.providers.ec2\_boto.BotoCloudProvider method), 42  
 is\_instance\_running() (elasticcluster.providers.gce.GoogleCloudProvider method), 43

## L

list\_instances() (elasticcluster.providers.gce.GoogleCloudProvider method), 43  
 ListClusters (class in elasticcluster.subcommands), 45  
 ListNodes (class in elasticcluster.subcommands), 45  
 ListTemplates (class in elasticcluster.subcommands), 45  
 load\_cluster() (elasticcluster.conf.Configurator method), 39

## M

MemRepository (class in elasticcluster.repository), 45

## N

Node (class in elasticcluster.cluster), 36

## P

pprint() (elasticcluster.cluster.Node method), 37  
 pre\_run() (elasticcluster.subcommands.AbstractCommand method), 45

## R

read\_config() (elasticcluster.conf.ConfigReader method), 37

remove\_node() (elasticcluster.cluster.Cluster method), 35  
 ResizeCluster (class in elasticcluster.subcommands), 46

## S

save\_or\_update() (elasticcluster.repository.AbstractClusterRepository method), 44  
 save\_or\_update() (elasticcluster.repository.ClusterRepository method), 45  
 save\_or\_update() (elasticcluster.repository.MemRepository method), 45  
 setup() (elasticcluster.cluster.Cluster method), 35  
 setup() (elasticcluster.subcommands.AbstractCommand method), 45  
 setup() (elasticcluster.subcommands.Stop method), 46  
 setup\_cluster() (elasticcluster.providers.AbstractSetupProvider method), 40  
 setup\_cluster() (elasticcluster.providers.ansible\_provider.AnsibleSetupProvider method), 41  
 SetupCluster (class in elasticcluster.subcommands), 46  
 SftpFrontend (class in elasticcluster.subcommands), 46  
 Singleton (class in elasticcluster.helpers), 39  
 SshFrontend (class in elasticcluster.subcommands), 46  
 Start (class in elasticcluster.subcommands), 46  
 start() (elasticcluster.cluster.Cluster method), 35  
 start() (elasticcluster.cluster.Node method), 37  
 start\_instance() (elasticcluster.providers.AbstractCloudProvider method), 39  
 start\_instance() (elasticcluster.providers.ec2\_boto.BotoCloudProvider method), 42  
 start\_instance() (elasticcluster.providers.gce.GoogleCloudProvider method), 43  
 startup\_timeout (elasticcluster.cluster.Cluster attribute), 36  
 Stop (class in elasticcluster.subcommands), 46  
 stop() (elasticcluster.cluster.Cluster method), 36  
 stop() (elasticcluster.cluster.Node method), 37  
 stop\_instance() (elasticcluster.providers.AbstractCloudProvider method), 40  
 stop\_instance() (elasticcluster.providers.ec2\_boto.BotoCloudProvider method), 42  
 stop\_instance() (elasticcluster.providers.gce.GoogleCloudProvider method), 44  
 update() (elasticcluster.cluster.Cluster method), 36



`update_ips()` (elasticcluster.cluster.Node method), [37](#)

## V

`validate()` (elasticcluster.conf.ConfigValidator method), [37](#)