

---

# Enos Documentation

*Release 2.2.0*

**discovery**

**Aug 16, 2017**



---

## Contents:

---

<b>1</b>	<b>Enos workflow</b>	<b>3</b>
1.1	Getting Started . . . . .	3
1.2	Provider . . . . .	4
1.3	Enos command line . . . . .	11
1.4	Benchmarks . . . . .	11
1.5	Customizations . . . . .	13
1.6	Network Emulation . . . . .	15
1.7	Analysis . . . . .	17
1.8	Contribute . . . . .	17
<b>2</b>	<b>Why Enos ?</b>	<b>21</b>
<b>3</b>	<b>License</b>	<b>23</b>
<b>4</b>	<b>Indices and tables</b>	<b>25</b>



**Hint:** The source code is available at <https://github.com/BeyondTheClouds/enos>

---

Enos deploys OpenStack and targets reproducible experiments. It allows easy:

- deployment of the system
- customization of the system
- benchmarking of the system
- visualization of various metrics

Enos is developed in the context of the [Discovery](#) initiative.



A typical experiment using Enos is the sequence of several phases:

- `enos up` : Enos will read the configuration file, get machines from the resource provider and will prepare the next phase
- `enos os` : Enos will deploy OpenStack on the machines. This phase rely highly on Kolla deployment.
- `enos init-os` : Enos will bootstrap the OpenStack installation (default quotas, security rules, ...)
- `enos bench` : Enos will run a list of benchmarks. Enos support Rally and Shaker benchmarks.
- `enos backup` : Enos will backup metrics gathered, logs and configuration files from the experiment.

## Getting Started

### Installation

```
$ pip install enos
```

You may prefer to go with a virtualenv. Please refer to the [virtualenv](#) documentation and the rest of this section for further information.

If virtualenv is missing:

```
$ pip install virtualenv --user      # Install virtualenv
$ export PATH=~/.local/bin/:${PATH} # Put it into your path
```

Then install enos inside a virtualenv:

```
$ mkdir my-experiment && cd my-experiment
$ virtualenv venv
$ source venv/bin/activate
(venv) $ pip install enos
```

**Note:** The latest *packaged* version of enos will install the latest *stable* version of OpenStack. If you want to install the development version of OpenStack, you should install enos from sources (see [Contribute](#)).

---

## Configuration

To get started you can get the sample configuration file and edit it:

```
$ curl https://raw.githubusercontent.com/BeyondTheClouds/enos/stable/ocata/  
↪reservation.yaml.sample --output reservation.yaml  
$ <editor> reservation.yaml
```

The configuration may vary from one provider to another, please refer to the dedicated *Provider* configuration

**Note:** If a key is defined several times in the configuration file, only the last occurrence will be taken into account. In particular to switch from one provider to another, you can move down the key `provider` and its associated `resources` key.

---

## Deployment

Once your configuration is done, you can launch the deployment :

```
(venv) $ enos deploy
```

## Provider

Enos offers to test different OpenStack deployments over some resources. In the context of Enos, a resource is anything Enos can SSH on and start a Docker daemon. Thus, a resource could be a bare-metal machine, a virtual machine, or a container resource depending on the testbed used for conduction the experiments. To get these resources, Enos relies on a notion of *provider* and already implements the followings:

### Static

The static provider reuses already available resources (machines, network) to deploy OpenStack on.

### Installation

Refer to the *Getting Started* section to install Enos.

### Configuration

The static provider requires already running resources to deploy OpenStack on. Information in the provider description tells Enos where these resources are and how to access to them. Concretely, you have to fill the following information to properly configure Enos with already running resources.



```

provider: # Configuration for the static provider
          # All keys are mandatory
type: static
network:
  extra_ips: # Array of (virtual) IPs to be assigned
             # during the deployment (e.g HAProxy)
             # At least 5 IPs routable through the
             # APIs interface
             # Openstack external network will be configured according
             # to the following keys/values. Ips have to be routable through
             # the external interface
  start: # First available ip
  end: # Last available ip
  cidr: # cidr notation to describe the pool of ips
  gateway: # Ip address of the gateway
  dns: # Ip address of the DNS
eths:
  - eth1 # Name of the APIs interface
  - eth2 # Name of the external interface

resources: # An object with roles and there associated resources
control: # The role control, followed by resource info:
  address: # Ip of the control node
  alias: # Name for the node (optional)
  user: # ssh user name (optional)
  keyfile: # path to ssh private key (optional)
  port: # ssh port (optional)
compute: # A role can also have a list of resources
  - address: # Ip address of the first compute
    alias: # ...
    user: # ...
    keyfile: # ...
    port: # ...
  - address: # Ip address of the second compute
  - address: # Ip address of the third compute

```

## Grid'5000

### Installation

Connect to the frontend of your choice. Then, refer to the *Getting Started* section to install Enos.

### Configuration

The provider relies on cluster names to group the wanted resources. For example the following is a valid resources description :

```

provider: g5k

resources:
  paravance:
    control: 1
    network: 1
  parsilo:
    compute: 10

```

### Deployment

We suggest you to run the deployment from a dedicated node (especially for large deployment). For now the recommended way to do so is to reserve one node prior of your reservation. In the case of an interactive deployment:

```
frontend) oarsub -I -l 'walltime=2:00:00'
node) source venv/bin/activate
node) <edit configuration>
node) enos deploy
```

Enos will need to access the G5K api from this node. In order to do so you need to instruct execo to use the external endpoint of the API :

```
cat ~/.execo.conf.py:
g5k_configuration = {
    'api_username': '<your login>',
    'api_uri': "https://api-ext.grid5000.fr/3.0/",
}
```

You will be asked for your password at start.

### Default provider configuration

The provider comes with the following default options :

```
provider:
  type: g5k
  name: Enos
  walltime: 02:00:00
  env_name: jessie-x64-min
  reservation: None
  vlans: '{rennes: "{type=kavlan}/vlan=1"}'
  role_distribution: strict
  single_interface: false
  user: root
```

They can be overridden in the configuration file.

## Vagrant

### Installation

To get started with the vagrant provider, you need to install

- [Vagrant](#)

Then, refer to the *Getting Started* section to install Enos.

### Configuration

The provider relies on virtual machine sizes to group the wanted resources. For example the following is a valid resources description:

```
provider: vagrant

resources:
  medium:
    control: 1
    network: 1
  small:
    compute: 1
```

The list of the sizes may be found [here](#).

By default virtualbox will be used. See below to learn how to change the default virtualbox backend.

### Use libvirt as the backend for Vagrant

Declaring your provider options as the following will spin up virtual machines using libvirt. Note that `vagrant libvirt` must be installed on your system.

```
provider:
  type: vagrant
  backend: libvirt
```

### Default provider configuration

The provider comes with the following default options:

```
provider:
  type: vagrant
  backend: virtualbox
  box: debian/jessie64
  user: root
```

They can be overridden in the configuration file.

## Openstack

The OpenStack provider allows you to use Enos on an OpenStack cloud. In other words this lets you run OpenStack on OpenStack. In the following, the under-cloud is the underlying OpenStack infrastructure, the over-cloud is the OpenStack configured by Enos.

The over-cloud configured by Enos needs a set of resources to be present on the under-cloud. The first step in the deployment workflow consists in checking or creating such resources. Some resources are mandatory and must be present before the deployment (base image, keypairs, ...), some others can be created or reused during the deployment (private networks). For the latter, you can use the default values set by the provider.

For specific under-clouds (e.g Chameleon), specific providers deriving from the OpenStack provider may be used. They will enforce more default values that fit the under-cloud specificities (e.g specific DNS, base image, ...)

### Installation

Please refer to *Getting Started*.

## Configuration

The provider relies on flavor names to group the wanted resources. For example the following is probably a valid resources description.

```
provider:
  type: openstack
  <options see below>

resources:
  ml.medium:
    control: 1
    network: 1
  ml.small:
    compute: 10
```

## Default provider configuration

The OpenStack provider is shipped with the following default options. These options will be set automatically and thus may be omitted in the configuration file.

```
provider:
  type: openstack
  # True if Enos needs to create a dedicated network to work with
  # False means you already have a network, subnet and router to
  # your ext_net configured
  configure_network: true

  # Name of the network to use or to create
  # It will be use as external network for the upper-cloud
  network: {'name': 'enos-network'}

  # Name of the subnet to use or to create
  subnet: {'name': 'enos-subnet', 'cidr': '10.87.23.0/24'}

  # DNS server to use when creating the network
  dns_nameservers: ['8.8.8.8', '8.8.4.4']

  # Floating ips pool
  allocation_pool: {'start': '10.87.23.10', 'end': '10.87.23.100'}

  # Whether one machine must act as gateway
  # - False means that you can connect directly to all the machines
  # started by Enos
  # - True means that one machine will be assigned a floating ip and used
  # as gateway to the others
  gateway: true
```

These options can be overridden in the provider config.

## Mandatory provider configuration

The following parameters must be present in your configuration file

```

provider:
  type: openstack
  key_name:      # the ssh keypair name to use
  image_name:    # base image to use to boot the machines from
                 # debian8 or ubuntu16.04
  user:         # user to use when connecting to the machines
  network_interface: # network interface of the deployed machines

```

## Deployment

Enos will interact with the remote OpenStack APIs. In order to get authenticated you must source your rc file. To use Enos on Openstack there are two distinct cases :

- If you have direct access to all your machines (`gateway: false`), you can launch the deployment with :

```
enos deploy
```

---

**Hint:** In this case, prior to the Enos deployment, you have probably started a machine to act as a frontend. This machine is in the same network as those used by Enos

---

- If you don't have direct access to all your machines (`gateway: true`)

*init* phase rely on accessing the OpenStack APIs of the over Cloud. If you choose to deploy from your local machine, those APIs are probably unreachable. Check the floating ip assigned to one of your machine and create a proxy socks.

```

# terminal 1
enos up
enos os

# terminal 2
ssh -ND 2100 user@<floating-ip>

# terminal 1
export http_proxy=socks5://127.0.0.1:2100
pip install requests[socks] # install requests support to socks
enos init

```

Note that the proxy socks allows you to use any *openstack* command directly to the over-cloud.

## Chameleon Cloud (Bare Metal)

This provider is an OpenStack based provider where some options are set to fit the following platforms :

- <https://chi.uc.chameleoncloud.org/>
- <https://chi.tacc.chameleoncloud.org/>

## Deployment

You need to install *python-blazarclient* to interact with the lease system of Chameleon :

```
pip install git+https://github.com/openstack/python-blazarclient
```

### Configuration

As more default values can be enforced automatically, the following is a valid resources description.

```
provider:
  type: chameleonbaremetal
  key_name: 'enos-key' # must be present prior to the execution

resources:
  storage: # use "storage" machine type for these roles
    control: 1
    network: 1
  compute: # use "compute" machine type for these roles
    compute: 10
```

Note that on Chameleon, they are two groups of machines : compute and storage.

### Default provider configuration

The following options will be set automatically and thus may be omitted in the configuration file.

```
provider:
  type: chameleonbaremetal
  # Name of the Blazar lease to use
  lease_name: enos-lease
  image_name: CC-Ubuntu16.04
  user: cc
  configure_network: False
  network: {name: sharednet1}
  subnet: {name: sharednet1-subnet}
  dns_nameservers: [130.202.101.6, 130.202.101.37]
  # Name of the network interface available on the nodes
  network_interface: eno1
  # Experiment duration
  walltime: "02:00:00"
```

These options can be overridden in the provider config.

On <https://chi.tacc.chameleoncloud.org/> the subnet must be `subnet: {'name': 'shared-subnet1'}`

**Warning:** A shared-network is used and may limit the features of the over-cloud (e.g floating ips)

### Chameleon Cloud (KVM)

This provider is an OpenStack based provider where some options are set to fit the following platform :

- <https://openstack.tacc.chameleoncloud.org>

### Configuration

As more default values can be enforced automatically, the following is a valid resources description.

```

provider:
  type: chameleonkvm
  key_name: 'enos-key' # must be present prior to the execution

resources:
  ml.large:
    control: 1
    network: 1
  ml.medium:
    compute: 10

```

## Default provider configuration

The following options will be set automatically and thus may be omitted in the configuration file :

```

provider:
  type: chameleonkvm
  image_name: CC-Ubuntu16.04
  user: cc
  dns_nameservers: [129.114.97.1, 129.114.97.2, 129.116.84.203]
  network_interface: ens3

```

These options can be overridden in the provider config.

## Custom Provider

See *Write a new provider*.

## Enos command line

Once installed Enos give you access to its command line. Please refer to the output of `enos -h`. For a specific command you can use `enos <command> -h`

## Benchmarks

Benchmarks are run by Enos by the mean of a workload description. A workload is a set of scenarios grouped by type. A workload is launched with the following command:

```
(venv) $ enos bench --workload=workload
```

enos will look into the `workload` directory for a file named `run.yml`. This file is the description of the workload to launch. One example is given below:

```

rally:
  enabled: true # default is true
  args:
    concurrency:
      - 1
      - 2
      - 4

```

```
times:
  - 100
scenarios:
  - name: boot and list servers
    enabled: true # default is true
    file: nova-boot-list-cc.yml
    args:
      sla_max_avg_duration: 30
      times: 50
```

This will launch all the scenarios described under the `scenarios` keys with all the possible parameters. The parameters are calculated using the cartesian product of the parameters given under the `args` keys. Locally defined args (scenario level) shadow globally defined args (top level). The same mechanism is applied to the `enabled` values. The scenario must be parameterized accordingly. The key (`rally` here) defines the type of benchmark to launch: in the future we may support other type of scenarios.

After running the workload, a backup of the environment can be done through `enos backup`.

## Rally

Enos supports natively Rally scenarios. Please refer to the Rally documentation for any further information on this benchmarking tool.

Supported keys :

- `name`: the name of the scenario. Can be any string.
- `file`: must be the path to the scenario file. The path is relative to the `workload` directory
- `enabled`: Whether to run this scenario
- `args`: Any parameters that can be understood by the rally scenario
- `plugin`: must be the path to the plugin. The path is relative to the `workload` directory

## Shaker

Enos supports natively Shaker scenarios. Please refer to the Shaker documentation for any further information on this benchmarking tool.

Supported keys :

- `name`: the name of the scenario. Can be any string.
- `file`: must be the alias of the scenario. Enos don't support custom scenario yet.
- `enabled`: Whether to run this scenario

## Osprofiler

Supporting OSProfiler in Rally benchmarks is planned for Q3 2017.



## Customizations

### Changing Kolla / Ansible variables

Custom Kolla / Ansible parameters can be put in the configuration file under the key `kolla`. The complete list of Kolla variables can be found [here](#).

For instance, Kolla uses the `openstack_release` parameter to fix the OpenStack version to deploy. So, Enos tells Kolla to deploy the 4.0.0 version with:

```
kolla:
  openstack_release: "4.0.0"
```

Note that the Kolla code varies from one version of OpenStack to another. You should always target a version of Kolla code that support the deployment of the expected OpenStack. To do so, you can change the git repository/reference of Kolla code with:

```
kolla_repo: "https://git.openstack.org/openstack/kolla-ansible"
kolla_ref: "stable/ocata"
```

You can also your own local clone of kolla-ansible with:

```
kolla_repo: "file:///path/to/local/kolla-ansible"
```

#### Note on the network interfaces:

Providers do their best to configure the network decently. This probably doesn't cover all the possible use cases. But, if you know what interfaces are configured by the provider you can specify a more precise allocation under the `kolla` key. For instance:

```
kolla:
  network_interface: eth1
  neutron_external_interface: eth2
  tunnel_interface: eth3
```

### Running from kolla/master

if you want to live on the bleeding edge you can run the latest Kolla code with the latest built kolla images.

```
kolla_repo: "https://git.openstack.org/openstack/kolla-ansible"
kolla_ref: "master"

kolla:
  docker_namespace: "beyondtheclouds"
  openstack_release: "latest"
```

### Changing the topology

Let's assume you want to run the `nova-conductor` in a dedicated node:

1. Add a new node reservation in the configuration file:

```
paravance:
  control: 1
  network: 1
  compute: 1
  conductor-node: 1
```

2. Create a new inventory file in the `inventories` subdirectory (copy paste the sample inventory) and change the group of the conductor service:

```
[nova-conductor:children]
conductor-node
```

3. In the configuration file, points the inventory to use to this new inventory.

4. Launch the deployment as usual, and you'll get the `nova-conductor` on a dedicated node.

## Configuration tuning

At some point, Kolla default parameters won't fit your needs. Kolla provides a mechanism to override custom section of configuration files but isn't applicable in our case (at least in the corresponding branch). So we implement a *quick and dirty* way of patching Kolla code to enable custom configuration files to be used (and by extension custom kolla code). See the possible patch declaration in `ansible/group_vars/all.yml`. Patches should be added in the configuration file of the experiment and you can rely on the `{{ cwd }}` key to link patches in your current working directory.

For instance, adding the following in the configuration file tells enos to look into the current working directory for a file called `mariadb_bootstrap.yml` that will replace the `kolla-ansible mariadb start` playbook.

```
patches:
- name: Patch mariadb start
  src: "{{ cwd }}/mariadb_bootstrap.yml
  dst: kolla/ansible/roles/mariadb/tasks/start.yml
  enabled: "yes"
```

## Ansible configuration

By default, Enos loads its own `ansible.cfg`. To use another Ansible configuration file, the `ANSIBLE_CONFIG` environment variable can be used. Further information can be found : [see here](#).

## Using a persistent registry with Ceph

Enos deploys a fresh registry that acts as a private docker registry mirroring the official one and cache containers close to your deployment resources.

To get a persistent registry you can use a persistent Ceph Rados Block Device for the registry backend. Image will be cached during the first deployment and reused for the subsequent deployments.

The relevant configuration section looks like this in your `reservation.yml`:

```
registry:
  ceph: true|false
  ceph_keyring: path to your keyring
  ceph_id: your ceph id
```

```
ceph_rbd: rbd in the form "pool/rbd"
ceph_mon_host: list of ceph monitor addresses
```

- `ceph: false` starts a fresh registry that caches the images for the duration of the experiment.
- `ceph: true` uses a registry whose backend is the existing `ceph_rbd` Ceph Rados Block Device at destination `ceph_mon_host` with the pool `ceph_id` and key `ceph_keyring`.

**Note:** The `reservation.yaml.sample` file provides an example of Ceph configuration that relies on the G5k Ceph of Rennes. [The G5k Ceph tutorial](#) will guide you to create your own Rados Block Device.

## Using a local registry

By default, Enos deploys a cache registry in the control node. You can tell enos to use a locally deployed insecure registry, that is accessible on port 4000, with :

```
registry:
  ip: my_ip
```

**Note:** If using a local registry, you can remove the `disco/registry` entry from the inventory, to avoid deploying the cache registry.

## Network Emulation

### Links description

Enos allows to enforce network emulation in terms of latency and bandwidth limitations.

Network constraints (latency/bandwidth limitations) are enabled by the use of groups of nodes. Resources *must* be described using a topology description instead of a resources description. The following example will define 4 groups named `grp1`, `grp2`, `grp3` and `grp4` respectively:

```
topology:
  grp1:
    paravance:
      control: 1
      network: 1
  grp[2-4]:
    paravance:
      compute: 1
```

Constraints are then described under the `network_constraints` key in the configuration file:

```
network_constraints:
  enable: true
  default_delay: 25ms
  default_rate: 100mbit
  default_loss: 0.1%
  constraints:
    - src: grp1
```

```
dst: grp[2-4]
delay: 10ms
rate: 1gbit
loss: 0%
symetric: true
```

To enforce the constraints, you can invoke:

```
enos tc
```

Note that The machines must be available, thus the *up* phase must have been called before.

As a result

- the network delay between every machines of `grp1` and the machines of the other groups will be 20ms (2x10ms: symetric)
- the bandwidth between every machines of `grp1` and the machines of the other groups will be 1 Gbit/s
- the packet loss percentage between every machines of `grp1` and the machines of the other groups will be 0%.
- the network delay between every machines of `grp2` and `grp3` (resp. `grp2` and `grp4`) (resp. `grp3` and `grp4`) will be 50ms
- the bandwidth between every machines of `grp2` and `grp3` (resp. `grp2` and `grp4`) (resp. `grp3` and `grp4`) will be 100Mbit/s
- the packet loss percentage between every machines of `grp2` and `grp3` (resp. `grp2` and `grp4`) (resp. `grp3` and `grp4`) will be 0.1%

## Checking the constraints

Invoking

```
enos tc --test
```

will generate various reports to validate the constraints. They are based on `fping` and `flent` latency and bandwidth measurements respectively. The reports will be located in the result directory.

## Notes

- `default_delay`, `default_rate`, `default_loss` are mandatory
- To disable the network constraints you can specify `enable: false` under the `network_constraints` key and launch again `enos tc`
- To exclude a group from any tc rule, you can add an optionnal `except` key to the `network_constraints`:

```
network_constraints:
  enable: true
  default_delay: 25ms
  default_rate: 100mbit
  default_loss: 0%
  constraints:
    - src: grp[1-3]
      dst: grp[4-6]
      delay: 10ms
      rate: 1gbit
```

```
symetric: true
except:
  - grp1
```

## Analysis

### Real-time

Setting `enable_monitoring: true` in the configuration file will deploy a monitoring stack composed of:

- Cadvisor and Collectd agents
- InfluxDB for metrics collection
- Grafana for the visualization / exploration

All these services are accessible on their default ports. For instance you'll be able to access grafana dashboards on port 3000 of the node hosting grafana.

Some dashboards are available [here](#).

### Post-mortem

TODO

### Annotations

Enos embeds an Ansible plugin to add annotations in Grafana. These annotations are marked points which provide rich information about events when hovered over. Enos uses the `ansible.cfg` file that loads the plugin. The plugin can be disabled by editing the line `callback_whitelist = influxdb_events` in the `ansible.cfg`. Note also that the plugin is automatically disabled when the monitoring tools are not deployed (i.e. when `enable_monitoring = false` is set in the configuration file).

Once the deployment is finished, a compatible dashboard must be used in Grafana to display annotations. An example of such dashboard is available [here](#).

## Contribute

All contributions are welcome on [BeyondTheClouds/enos](#). For any questions, feature requests, issues please use the [GitHub issue tracker](#).

### Install from sources and make them editable

```
$ git clone https://github.com/BeyondTheClouds/enos.git
$ cd enos
$ virtualenv venv
$ source venv/bin/activate
(venv) $ pip install -e .
```

## Get tox

```
(venv) $ pip install tox
```

## Running the tests

```
(venv) $ tox
```

## Running syntax checker

```
(venv) $ tox -e pep8
```

## Generate the documentation

```
(venv) $ tox -e docs
```

## Other Topics

### Write a new provider

The actual implementation gives providers for *Static* resources, *Vagrant*, *Grid'5000* and *Openstack* itself. If you want to support another testbed, then implementing a new provider is easy as 500 lines of Python code.

The new provider should follow the `provider.py` interface which consists in three methods: `init`, `destroy` and `default_config`. Another good starting point is the simple [static implementation](#).

### Init Method

The `init` method provides resources and provisions the environment. To let the provider knows what kind and how many resources should be provided, the method is fed with the `config` object that maps the reservations file. So a provider can access the resource description with:

```
rsc = config['resources']  
  
# Use rsc to book resources ...
```

At the end of the `init`, the provider should return a list of `host.py` that Enos can SSH on, together with a pool of available IP for OpenStack Network.

### Destroy Method

The `destroy` method destroys resources that have been used for the deployment. The provider can rely on the environment variable to get information related to its deployment.

## Default Provider Configuration Methods

The `default_config` specifies keys used to configure the provider with a `dict`. A key could be *optional* and so should be provided with a default value, or *required* and so should be set to `None`. The user then can override these keys in the reservation file, under the `provider` key. Keys marked as `None` in the `default_config` are automatically tested for overriding in the reservation file.

## Provider Instantiation

Enos automatically instantiates a provider based on the name specified in the `reservation.yaml`. For instance, based on the following reservation file,

```
provider: "my-provider"
```

Enos seeks for a file called `my-provider.py` into `enos/provider` and instantiates its main class. But sometimes, the provider requires extra information for its initialisation. The good place for this information is to put it under the `provider` key. In this case, the provider name should be accessible throughout the `type` key:

```
provider:
  type: "my-provider"
  extra-var: ...
  ...
```

Then the provider can access `extra-var` with `config['provider']['extra-var']`. Supported extra information is documented into the provider documentation.





## CHAPTER 2

---

### Why Enos ?

---

[https://en.wikipedia.org/wiki/Enos\\_\(chimpanzee\)](https://en.wikipedia.org/wiki/Enos_(chimpanzee))



## CHAPTER 3

---

### License

---

Enos runs performance stress workloads on OpenStack for postmortem analysis. Copyright (C) 2016 Didier Iscovery

This program is free software: you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation, either version 3 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program. If not, see <<http://www.gnu.org/licenses/>>.



## CHAPTER 4

---

### Indices and tables

---

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