

Combining Ansible and Terraform for CI

better together love story based on OVN-CI project

Arie Bregman
abregman@redhat.com

Szymon Datko
sdatko@redhat.com



Red Hat

20th October 2020



Arie Bregman

- DevOps Engineer
- Gamer with all heart
- Open Source aficionado



Szymon Datko

- DevOps & local Bash wizard
- Open Source software lover
- Computer Graphics enthusiast



Part I

The tools

What is Ansible?



ANSIBLE

- Ansible is an IT automation tool.
- You can run various tasks with Ansible:
 - configure systems,
 - deploy software,
 - anything you run in your scripts! :-)
- Open Source and actively developed.
- Written in Python, easy to start.
- Widely used by companies and developers!

How Ansible works?

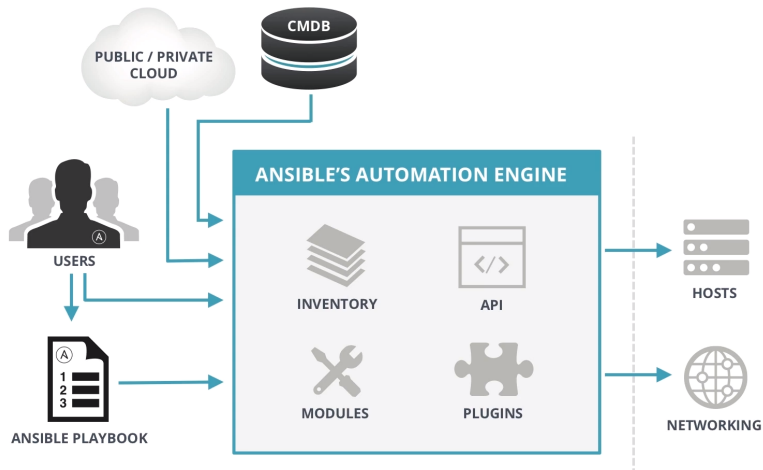


Image source: <https://www.ansible.com/resources/videos/quick-start-video>.

Example Ansible playbook

- Ansible playbooks are expressed in YAML format.
- As a user you specify what to run, where and under what conditions.
- Multiple playbooks can be grouped into a role.

```
1| # playbook.yml
2| ---
3| - name: Install DB servers
4|   hosts: databases
5|   remote_user: root
6|
7|   tasks:
8|     - name: Ensure postgresql is at the latest version
9|       package:
10|         name: postgresql
11|         state: latest
12|     - name: Ensure that postgresql service is started
13|       service:
14|         name: postgresql
15|         state: started
```

Running Ansible

- Run playbooks with: `ansible-playbook <playbook_file_path>`.
- Ad-hoc mode is also supported:
 - ① `ansible -m package -a "name:postgresql state:present",`
 - ② `ansible -m service -a "name:postgresql state:started".`
- Offline documentation for modules: `ansible-doc <module_name>`.

```
1| [user@db-server1 directory]$ ansible-playbook playbook.yml
2|
3| PLAY [Install DB servers] *****
4|
5| TASK [Ensure postgresql is at the latest version] *****
6| ok: [db-server1]
7|
8| TASK [Ensure that postgresql service is started] *****
9| ok: [db-server1]
10|
11| PLAY RECAP *****
12| db-server1          : ok=3   changed=0   unreachable=0   failed=0   skipped=0
```

What is Terraform?



- Use Infrastructure as Code to provision and manage any cloud, infrastructure or service.
- Terraform supports multiple providers:
 - clouds: OpenStack, GCP, AWS, Azure, ...;
 - containers: Docker, Kubernetes,
- Capable of resolving dependencies in resources.
- Aims to unify a way of resources management.
- Open Source, written in Go language.

How Terraform works?

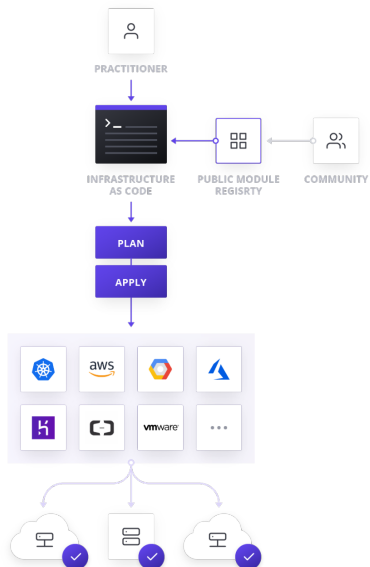


Image source: <https://www.terraform.io>.

Example Terraform manifest

- Terraform has its own declarative, configuration language,
 - HashiCorp Configuration Language (HCL).
- You specify which provider to use and the resources to manage.
- Configuration can be splitted to multiple files in the same directory.

```
1| provider "openstack" {
2|   auth_url      = "http://myauthurl:5000/v2.0"
3|   region       = "RegionOne"
4|   tenant_name  = "admin"
5|   user_name    = "admin"
6|   # password   = "" # If not given here, it will be read from env variable.
7| }
8|
9| resource "openstack_compute_instance_v2" "test-server" {
10|  name = "my test server"
11|  image_name = "centos8"
12|  flavor_name = "m1.medium"
13|  key_pair = "default-access-key"
14|  network {
15|    name = "ext-net"
16|  }
17| }
```

Running Terraform

- Run the following commands in your shell:

- terraform init
- terraform apply
- terraform destroy

- The `.terraform/` directory and `terraform.tfstate` file will be created.

```
1| [user@localhost directory]$ terraform init
2|
3| Initializing the backend...
4|
5| Initializing provider plugins...
6| - Finding latest version of terraform-providers/openstack
7| - Installing terraform-providers/openstack v1.32.0...
8| - Installed terraform-providers/openstack v1.32.0
9|
10| Terraform has been successfully initialized!
11| (...)
```



```
1| [user@localhost directory]$ terraform apply
2| (...)
3| Terraform will perform the following actions:
4|
5| # openstack_compute_instance_v2.test-server will be created
6| + resource "openstack_compute_instance_v2" "test-server" {
7|   + access_ip_v4      = (known after apply)
8|   + access_ip_v6      = (known after apply)
9|   + all_metadata      = (known after apply)
10|  + all_tags           = (known after apply)
11|  + availability_zone  = (known after apply)
12|  + flavor_id          = (known after apply)
13|  + flavor_name        = "m1.medium"
14|  + force_delete       = false
15|  + id                 = (known after apply)
16|  + image_id           = (known after apply)
17|  + image_name         = "centos8"
18|  + key_pair           = "default-access-key"
19|  + name               = "my test server"
20|  (...)
```



HashiCorp

Terraform



Part II

The combining

Why to combine Ansible and Terraform?

- "One tool to do everything" is not a goal by itself.
 - Though, for simple tasks you may not need to involve new toolkit.
- Recognizing the benefits of each tool will repay in complex setups.
 - Ansible is great for configuration management.
 - Terraform is great for provisioning resources.
- In the CI-related tasks: how do you track resources to clean up?



Image source: <https://knowyourmeme.com/photos/882534-swiss-army-knives>.

How to combine Ansible and Terraform?

- 1 In your shell: call Terraform, then call Ansible ;-)
 - Straightforward, yet elegant and comply with Unix philosophy.
 - Terraform can produce the Ansible inventory file for convenience.

```
$ terraform init
$ terraform apply -auto-approve
$ ansible-playbook --inventory terraform-hosts.ini
$ terraform destroy -auto-approve
```
- 2 Call Ansible and utilize Ansible module that runs Terraform.
 - The facts dictionary from Terraform output may be registered.
- 3 Call Terraform and use its provisioners mechanism to invoke Ansible.
 - *"Provisioners should only be used as a last resort. For most common situations there are better alternatives."* – Terraform documentation.
 - Details: <https://www.terraform.io/docs/provisioners/>.

Generating Ansible inventory with Terraform – configuration

- The `local_file` resource can be used to generate the inventory file.
- In the example below we create two servers with different names.

```
1| provider "openstack" {} // Cloud configuration will be read from env variables.
2|
3| resource "openstack_compute_instance_v2" "servers" {
4|   count = 2
5|
6|   name = "my test server ${count.index + 1}"
7|   image_name = "centos8"
8|   flavor_name = "m1.medium"
9|   key_pair = "default-access-key"
10|  network {
11|    name = "ext-net"
12|  }
13| }
14|
15| resource "local_file" "inventory" {
16|   content = templatefile("terraform-hosts.tpl", {
17|     servers = openstack_compute_instance_v2.servers[*]
18|   })
19|   filename = "terraform-hosts.ini"
20|   file_permission = "0644"
21| }
```



Generating Ansible inventory with Terraform – result

- The used template file has a syntax similar to popular **Jinja2**.
- Note that the outcome file is regular Terraform resource, so it will be:
 - created after terraform apply,
 - removed after terraform destroy.

Template file:

```
1| localhost  ansible_connection=local ansible_python_interpreter=auto_silent
2|
3| [nodes]
4| {% for index, server in servers ~}
5| node${index + 1} ansible_host=${server.access_ip_v4}
6| {% endfor ~}
```

Outcome inventory file:

```
1| localhost  ansible_connection=local ansible_python_interpreter=auto_silent
2|
3| [nodes]
4| node1  ansible_host=192.168.1.2
5| node2  ansible_host=192.168.1.3
```


Using Ansible module for running Terraform

Add to the Terraform configuration:

```
1| output "addresses" {
2|   value = openstack_compute_instance_v2.servers[*].access_ip_v4
3| }
```

Example Ansible playbook:

```
1| - name: Terraform example
2|   hosts: localhost
3|   tasks:
4|     - name: Create resources
5|       terraform:
6|         project_path: "{{ playbook_dir }}"
7|         force_init: true
8|         state: present
9|         register: terraform_run
10|     - debug:
11|       msg: "{{ terraform_run.outputs.addresses.value }}"
12|       # ok: [localhost] => {
13|       #   "msg": [
14|       #     "192.168.1.2",
15|       #     "192.168.1.3"
16|       #   ]
17|       #}
18|     - name: Cleanup resources
19|       terraform:
20|         project_path: "{{ playbook_dir }}"
21|         state: absent
```

Part III

Tips and pitfalls to avoid

Wait for connection and facts gathering

- Facts gathering fetches the useful information about the target host,
 - e.g. operating system details, file system layout, networking setup, etc.
- When used in the play level, it is performed before the tasks.
 - It will immediately try to connect and fail if host is not up and ready.

This is bad

```
1| - name: Incorrect play
2| gather_facts: true
3| hosts: my_nodes
4| tasks:
5|   - name: Ensure system is reachable
6|     wait_for_connection:
7|       delay: 5
8|       sleep: 5
9|       timeout: 300
10|
11|   - name: Install Java
12|     include_role:
13|       name: geerlingguy.java
```

This is good

```
1| - name: Correct play
2| gather_facts: false
3| hosts: my_nodes
4| tasks:
5|   - name: Ensure system is reachable
6|     wait_for_connection:
7|       delay: 5
8|       sleep: 5
9|       timeout: 300
10|
11|   - name: Gather facts
12|     setup:
13|
14|   - name: Install Java
15|     include_role:
16|       name: geerlingguy.java
```

Real-time output streaming from Ansible

- It may be useful for long-running tasks, like tests execution.
- No built-in solution so far, though there are proposals in community.
 - Details: <https://github.com/ansible/proposals/issues/92>.



Image source: <https://br.pinterest.com/pin/331366485055644550/>.

Real-time output streaming from Ansible – solution

Ideas:

- Patch Ansible, utilizing the callback plugin mechanism or so.
 - Example concept: <https://stackoverflow.com/a/54448411>.
- Do some Bash magic and SSH tunneling for this.
 - Will be really useful for **shell** tasks only.
 - The idea is to start a local process that will listen and display output.
 - Use the `--ssh-extra-args` Ansible argument to open the tunnel.



Diagram made with <https://draw.io/>.

Real-time output streaming from Ansible – implementation

In Bash script:

```
1| SOCAT_PID=-1
2| SOCAT_PORT=-1
3|
4| # selection of value for SOCAT_PORT goes here -- see next slide
5|
6| socat -t 3600 - "TCP-LISTEN:${SOCAT_PORT},fork,reuseaddr" &
7| SOCAT_PID=$!
8|
9| ansible-playbook --inventory hosts.ini \
10| --ssh-extra-args="-R 12345:localhost:${SOCAT_PORT}" \
11| "${BASE_DIR}/launch-unit-tests.yml"
```

In Ansible playbook:

```
1| - name: Example play
2|   hosts: node1
3|   tasks:
4|     ...
5|     - name: Long-running task
6|       shell:
7|         chdir: ovn
8|         cmd: |
9|           exec > /dev/tcp/localhost/12345 2>&1
10|          make
11|     ...
```



Real-time output streaming from Ansible – helpful addition

As a bonus, two helpful snippets for process cleanup and port selection.

In Bash script:


```
1| #
2| # Cleanup function for the script, trap ensures to execute it
3| #
4| function cleanup() {
5|     if [ "${SOCAT_PID}" -gt 0 ]; then
6|         kill "${SOCAT_PID}"
7|     fi
8| }
9| trap cleanup EXIT
10|
11| #
12| # Find random unused port
13| #
14| limit=20
15| for trial in $(seq 1 "${limit}"); do
16|     SOCAT_PORT=$(( ( ${RANDOM} % 57535 ) + 8000 ))
17|     if ! ( echo '' > "/dev/tcp/127.0.0.1/${SOCAT_PORT}" >/dev/null 2>&1 ); then
18|         break
19|     fi
20|     if [ "${trial}" -ge "${limit}" ]; then
21|         exit 1
22|     fi
23| done
```

Part IV

Final words

Summary – how we use Ansible and Terraform in OVN CI

- Project started to make OVN rock solid and comprehensively tested.
 - Recently OVN was announced to be default network driver for Neutron.
- Why Ansible and Terraform for CI?
 - Provides convenient managing and configuring environments.
 - Out-of-box allows to track usage and status of resources.
 - Less dependency on CI/automation system = easy to reuse by hand.



Thank you for your attention!

The slides are available: <http://datko.pl/ansible-terraform.pdf>



Combining Ansible and Terraform for CI

better together love story based on OVN-CI project

Arie Bregman
abregman@redhat.com

Szymon Datko
sdatko@redhat.com



Red Hat

20th October 2020