



Towards “Plug and Play” Research Cloud Node Deployments

Timothy Carr
University of Cape Town
Technical Specialist



cat /proc/timinfo

- Technical Specialist @ University of Cape Town for 17 years.
- Focused on the building and delivery of enterprise and research ICT applications to the community.
- Two demanding boys.
- Amazing supportive wife.



Agenda

- The current problem
- Possible solution
- What is IAC ?
- Provide me with the tools.
- Building consistency into my environment.
- Questions



The current problem

- Shared research user environments pose a few challenges
 - Unresponsive compute nodes
 - Non-existent or limited privilege escalation - (Docker, Reboot Node)
- Manually maintaining and scaling compute stacks.
- What about configuration drift.
- Manual operating system / application patch cycles.



Possible Solution

- Research Purpose Built Environments
 - Deployed and defined with IAC.
 - Self-maintained with dynamic updates, scale the infrastructure on demand.
 - Block storage on-demand (BeeGFS) or (CephFS)
 - Pre-defined or custom virtual machine images (private glance image repo)
 - Software no longer resides in the virtual machine images (light-weight deploy) but in containers ie: Docker or Singularity.
 - Allows for privilege escalation, researcher is able to reboot unresponsive nodes.



What is IAC ?

- IAC (infrastructure as code) is defined as just that, “building/deploying/scaling infrastructure with code” against a set of cloud API endpoints.
- These endpoints are public cloud endpoints from cloud providers such as GCP, AWS, Rackspace, Digital Ocean, iLifu.
- IAC can maintain small compute stacks initiated for a specific purpose to managing the deployment of entire data centers.



Provide me with the tools

Cloud Native Landscape

v0.9.9

Database & Data Analytics Streaming SCM Application Definition CI/CD

Scheduling & Orchestration Coordination & Service Discovery Service Management

Cloud-Native Storage Container Runtime Cloud-Native Network

Host Management / Tooling Infrastructure Automation Container Registries Secure Images Key Management

Platforms

Observability & Analysis

Public Private



github.com/cncf/landscape

This landscape is intended as a map through the previously uncharted terrain of cloud native technologies. There are many routes to deploying a cloud native application, with CNCF Projects representing a particularly well-traveled path.




CLOUD NATIVE
COMPUTING FOUNDATION



seriously..... provide me with some realistic tooling

- **Hashicorp Terraform.io:** IAC tool to maintain infrastructure.
- **Hashicorp Packer.io:** Build / maintain virtual machine images to be consistent / immutable.
- **Configuration Mgmt Tools:** Ansible, Puppet, Chef, Saltstack
- **Openstack Heat:** Similar to Terraform but with integration points to other Openstack services. Has the autoscaling templates built-in for reference to “Openstack Telemetry”.

Tool: Terraform.io



- Declarative configuration: User *declares* exactly which portion of the infrastructure needs to change.
- Execution Plan: A “plan” will display the output of what will be committed to the environment. Consider it something similar to a “dry-run”.
 - **“\$terraform plan -out plan.file”**
- Change Automation: Terraform drafts a resource graph which will display which components will be changed in which order.
- Create: Apply the infrastructure changes to the environment.
 - **“\$terraform apply plan.file”**
- Destroy: **“\$terraform destroy plan.file”**

Code Snips

```
1 # Configure the OpenStack Provider
2 provider "openstack" {
3   user_name = "username@uct.ac.za"
4   tenant_name = "Tim Project"
5   password = "#####"
6   auth_url = "https://api.ilifu.ac.za:5000/"
7   region = "UCT-ICTS-DC"
8 }
9 |
```

provider.tf

```
13 resource "openstack_compute_secgroup_v2" "worker-icmp" {
14   name = "worker-icmp"
15   description = "Worker ICMP"
16   rule {
17     from_port = -1
18     to_port = -1
19     ip_protocol = "icmp"
20     cidr = "10.0.0.0/16"
21   }
22 }
```

security.tf

```
86 # Request floating ip
87 resource "openstack_networking_floatingip_v2" "headnode" {
88   pool = "${var.external_network}"
89 }
90
91 # Attach floating ip to instance
92 resource "openstack_compute_floatingip_associate_v2" "headnode" {
93   floating_ip = "${openstack_networking_floatingip_v2.headnode.address}"
94   instance_id = "${openstack_compute_instance_v2.headnode.id}"
95 }
```

floating_ip.tf

instances.tf

variables.tf

```
25 variable "flavor_workernode" {
26   default = "ilifu-A"
27 }
28
29 variable "private_network" {
30   default = {
31     network = "private_network"
32     subnet_name = "private_subnet"
33     cidr = "10.0.0.0/16"
34   }
35 }
36
37 variable "worker_count" {
38   default = "5"
39 }
40 }
```

```
10 # Create Meerkat SLURM-headnode instance
11 resource "openstack_compute_instance_v2" "headnode" {
12   name = "${var.head_name}"
13   image_name = "${var.image}"
14   flavor_name = "${var.flavor_headnode}"
15   key_pair = "${openstack_compute_keypair_v2.user_key.name}"
16   security_groups = ["${openstack_compute_secgroup_v2.ssh.id}"]
17   user_data = "${file("scripts/deploy.sh")}"
18   network {
19     uuid = "${openstack_networking_network_v2.private_network.id}"
20   }
21 }
22 # Install system in volume
23 block_device {
24   volume_size = "${var.volume_headnode}"
25   destination_type = "volume"
26   delete_on_termination = true
27   source_type = "image"
28   uuid = "${data.openstack_images_image_v2.centos_current.id}"
29 }
```

```
30 # Create Meerkat worker instances
31 resource "openstack_compute_instance_v2" "worker" {
32   count = "${var.worker_count}"
33   name = "${format("${var.worker_name}-%02d", count.index+1)}"
34   image_name = "${var.image}"
35   flavor_name = "${var.flavor_workernode}"
36   key_pair = "${openstack_compute_keypair_v2.user_key.name}"
37   security_groups = ["${openstack_compute_secgroup_v2.ssh.id}",
38     "${openstack_compute_secgroup_v2.worker-icmp.id}"]
39   user_data = "${file("scripts/first-boot.sh")}"
40   network {
41     uuid = "${openstack_networking_network_v2.private_network.id}"
42   }
43 }
```



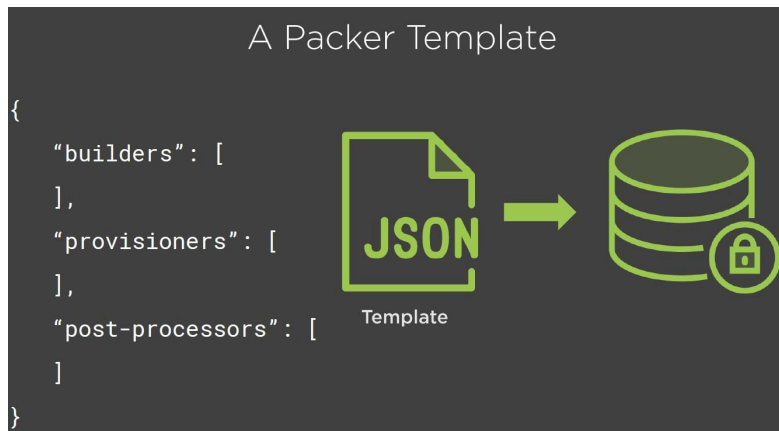
Installing Terraform

[#####100%]

- Terraform is written in GoLang.
- Single binary file and no dependency installs required
- Multiple Provider Support - AWS, GCP, Openstack, Vmware, many many others
- Information provided in your **provider.tf** downloads the relevant provider modules.
 - “\$ terraform init “
- Installation on your desktop machine or a shared environment.

Tool: Packer.io

- Written in GO / X-platform
- Build / Automate/ Maintain virtual machine images.
- Packer could be used as an alternative to configuration management.
- Packer includes support for provisioners.
- The packer JSON defines everything which is in the image.
- Packer automatically builds and ships to Openstack Glance.
- Additional metadata tags can be included about the image import.





Building consistency into my environment

- It is important to maintain a robust and consistent environment.
 - If it is broken, destroy and redeploy.
 - Your reference data should be elsewhere so your compute environment should be flexible to destroy.
 - The environment should be immutable.
 - Software should be built into containers.



Questions ?