Using Terraform and Ansible to provision a HPC Cluster

Rik Kisnah
rikesh@gmail.com

09-09-2021

Agenda

- Introduction
- Terraform and Ansible in the DevOps World
- HPC in the Cloud
- The love triangle – Terraform, Ansible and HPC
- Demo

Introduction

- Principal Engineer with more than ten years in Cloud Industry (Systems engineering/QA/Automation)
- *Disclaimer: I am here in a personal capacity. I don’t represent or speak behalf of current or previous employers*
- Based in Seattle
- Avid reader of books and listener to audio (I read about 2-4 books per month)
- Favorite programming language is Python

Terraform

- Terraform is a infrastructure as code
- A hashicorp product which supports multiple cloud vendors – Oracle, AWS, Azure etc
  - https://www.terraform.io/intro/index.html
- Example of a simple snippet to launch a host

```
resource "oci_core_instance" "ubuntu_instance" {
    # Required
    availability_domain = data.oci_identity_availability_domains.ads.availability_domains[0].name
    compartment_id = "<compartment-ocid>"
    shape = "VM.Standard.E2.1"
    source_details {
        source_id = "<source-ocid>"
        source_type = "image"
    }
    # Optional
    display_name = "<your-ubuntu-instance-name>"
    create_vnic_details {
        assign_public_ip = true
        subnet_id = "<subnet-ocid>"
    }
    metadata = {
        ssh_authorized_keys = file("<ssh-public-key-path>")
    }
    preserve_boot_volume = false
}
```

Ansible

- Open Source redhat project that automates provisioning without the use of external agents (it uses SSH)

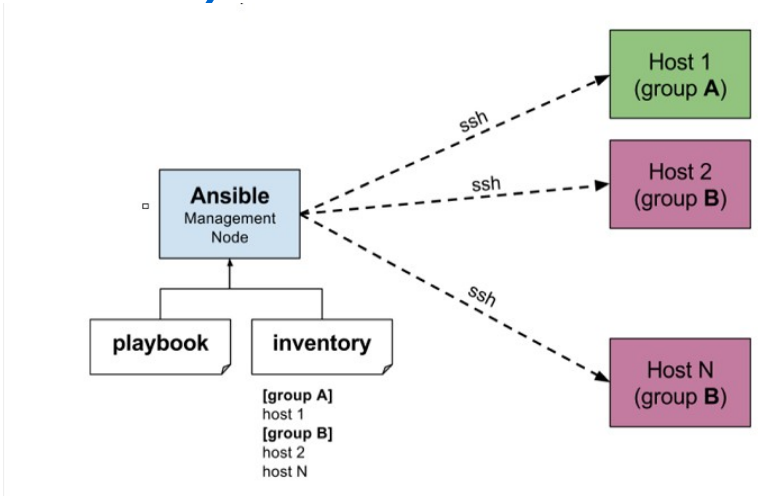

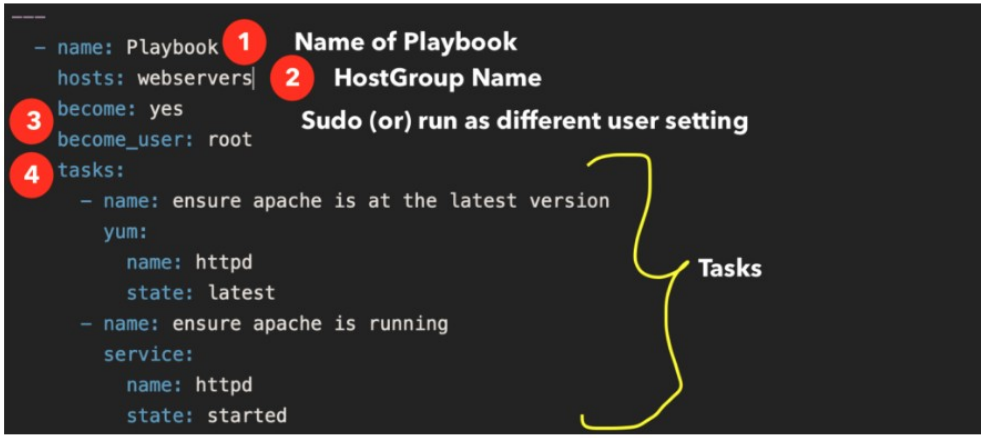

HPC in the cloud

- High Power Computing is ability to process huge amount of data in a short span of time (uses in AI/ML/Big Data industries)
- Vendors or on premise solutions configures hosts with special NICs (RDMA) or network (infiniband) to optimize latency
- HPC Cloud providers provides multiple implementations – Infiniband(Azure), ROCEv1, ROCEv2 (OCI) .Uses RDMA NICs for networking (Mellanox)
- Great whitepaper on HPC and RDMA

  https://www.open-mpi.org/papers/euro-pvmmpi-2006-hpc-protocols/euro-pvmmpi-2006-hpc-protocols.pdf

Example of HPC cloud provider

- https://www.oracle.com/cloud/hpc/
- https://docs.oracle.com/en/solutions/deploy-hpc-on-oci/index.html#GUID-FB8F3DC8-CA7B-47CE-85F6-426CE92B74A7
- Provides a clusters of HPC host with and without GPUs on fast RDMA network

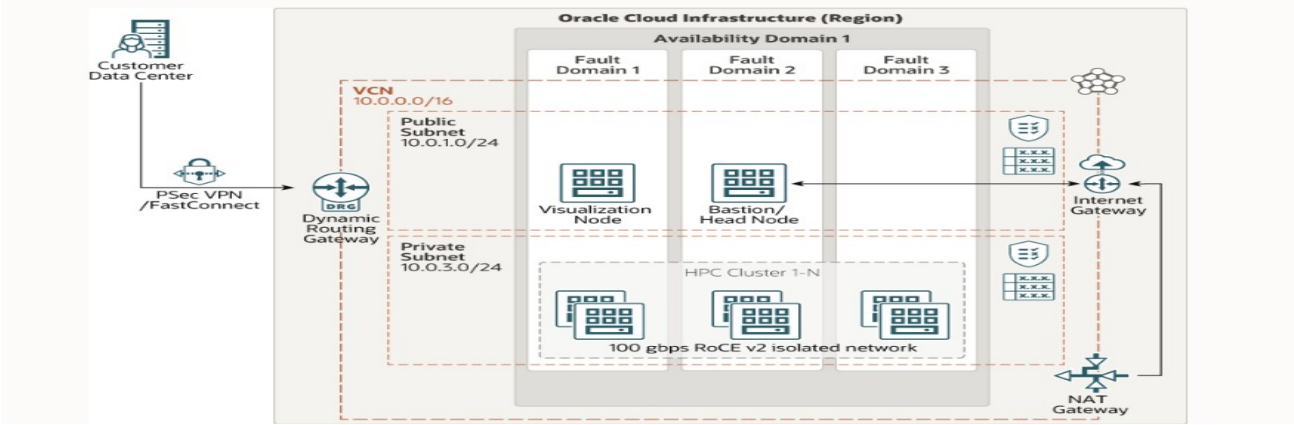

Terraform+Ansible=Cluster

- Use Terraform to launch the cluster and a bastion
- Use Terraform to set up ansible roles on the bastion
- Use roles to set up the cluster in the the network
- Set up packages / NFS / Slurm / NICs etc
- Open source solution: https://github.com/oci-hpc/oci-quickstart-hpc

Demo

- Show how to launch a cluster with the TF and Ansible
- Ping a host on the RDMA network

Q&A