

3 - Infrastructure Provisioning (IaC)

#	Step	
1	Terraform Cloud workspaces	<p>We need to maintain two workspaces - one for the Fabric Kubernetes cluster and one for the openIDL applications.</p> <p>To create the workspaces use the tool located in senofi/openidl-devops:</p> <ol style="list-style-type: none">Go to <code>openidl-devops/aws-infrastructure/environments/<env-folder>/terraform-cloud</code> and run <pre>terraform plan</pre> <p>If everything looks ok, execute <code>terraform apply</code>. This should create two workspaces and a var set in Terraform Cloud.</p> <ol style="list-style-type: none">The apply from step 1 should fail as the execution of the script is happening on the Terraform Cloud servers and the local config in <code>~/.terraformrc</code> is only useful to the communication from the localhost to the Terraform Cloud servers. In order to fix this, find in Terraform Cloud the new workspace that has the format of <code><org_id>-<env>-tf_cloud</code> and go to <code>Variables</code> section where a new environment variable will be needed with the name <code>TFE_CLOUD</code>, type environment variable and value the Terraform token that was previously saved in <code>~/.terraformrc</code>Create a new KMS key (symetric, encrypt/decrypt) in the AWS console. The name is not important but use a meaningful name that will associate it with this environment. Use it to populate the property in the next stepGo to <code>openidl-devops/automation/terraform-cloud</code> and update <code>configuration.properties</code> Make sure that the var set name matches what is in Terraform CloudCreate SSH keys <pre>ssh-keygen -t rsa -f app_eks_worker_nodes_ssh_key.pem ssh-keygen -t rsa -f blk_eks_worker_nodes_ssh_key.pem ssh-keygen -t rsa -f bastion_ssh_key.pem</pre> <ol style="list-style-type: none">Populate the variable set by executing the following command in <code>openidl-devops/automation/terraform-cloud</code> <pre>pip install -r requirements.txt python populate-variable-set.py</pre> <ol style="list-style-type: none">Copy the contents of the public keys and populate them in Terraform Cloud UI under Variable Sets <the newly created varset> <ol style="list-style-type: none"><code>app_eks_worker_nodes_ssh_key.pem.pub</code><code>bastion_ssh_key.pem.pub</code><code>blk_eks_worker_nodes_ssh_key.pem.pub</code>
2	Configure Jenkins	

1. Set Jenkins node label 'openidl' in Kubernetes Cloud by going to Manage Jenkins > Manage Nodes and Clouds > Configure Clouds. Make sure that under Pod Template details the labels field contains the value 'openidl'.

Pod Templates

Use of Pod templates are scheduled agents

Pod Template Name ?

default

Pod Template details ^ Edit

Namespace ?

jenkins

Labels ?

openidl

Usage ?

Use this node as much as possible

Pod template to inherit from ?

Also, remove the prepopulated 'sleep' command if it is set on the pod template:

☐ Always pull image ?

Working directory ?

/home/jenkins/agent

Command to run ?

Arguments to pass to the command ?

\$(computer.jnlpmac) \$(computer.name)

☐ Allocate pseudo-TTY ?

2. Create the Terraform Job Template

- a. Terraform Token Secret - Login to Jenkins go to Manage Jenkins > Manage Credentials > Stores scoped to Jenkins (Jenkins) > Global Credentials (unrestricted) > Add credentials

Kind

Secret text

Scope

Global (Jenkins, nodes, items, all child items, etc)

Secret

.....

ID

TF_BEARER_TOKEN

Description

Jenkins to authenticate TF

OK

- b. Choose Kind as secret text, enter secret text like Token in "secret" field and name the secret ID as unique since it will be used in pipeline code.
- c. Git Credentials - Add a new credential

		<p>3. Terraform Job</p> <ol style="list-style-type: none"> Go to Jenkins New Item. Use a name such as <code>Terraform Job</code> Select job type as PIPELINE and proceed. Select Definition as Pipeline Script from SCM Select SCM as Git Key in the Infrastructure code repository (openidl-gitops) URL. Select the Git credential created above Specify the relevant branch "refs/heads/<branch-name>". Set script path to <code>jenkins-jobs/jenkinsfile-tf</code>
3	Run Terraform Job	<ol style="list-style-type: none"> Run the Jenkins Terraform Job Open the console log for the job. Once the job asks for an input accept and choose the apply option The job runs a second plan into the Kubernetes workspace in Terraform Cloud. When asked - accept and apply the changes Go to the AWS Console and find EKS (Elastic Kubernetes Service). Choose the blk cluster and go to Add-Ons. Find the EBS plugin and add it to the list. The plugin makes sure volumes could be created in Kubernetes
4	Add Amazon EBS CSI Driver Add-on	<ol style="list-style-type: none"> Go to AWS Console Elastic Kubernetes Service Find the new app and k8s clusters that were created in the previous step Open each of them and go to Add-ons tab Select Get more add-ons and select the Amazon EBS CSI Driver Choose Next on the bottom of the page Review and click Create button Repeat for the other cluster