



The Splunk Derp Gun

Automating Splunk deployments in the Cloud
with Ansible and Terraform

Robert Johansson – Timothy Mahoney
Securelink Sweden AB

About us:

- SecureLink Malmö
- Managed SIEM Team
- Book release 2020/04/01



A book filled with good intentions.



Cloudy With a Chance of Logs.

Broken Pipe Edition

O RLY?

Robert & Timothy

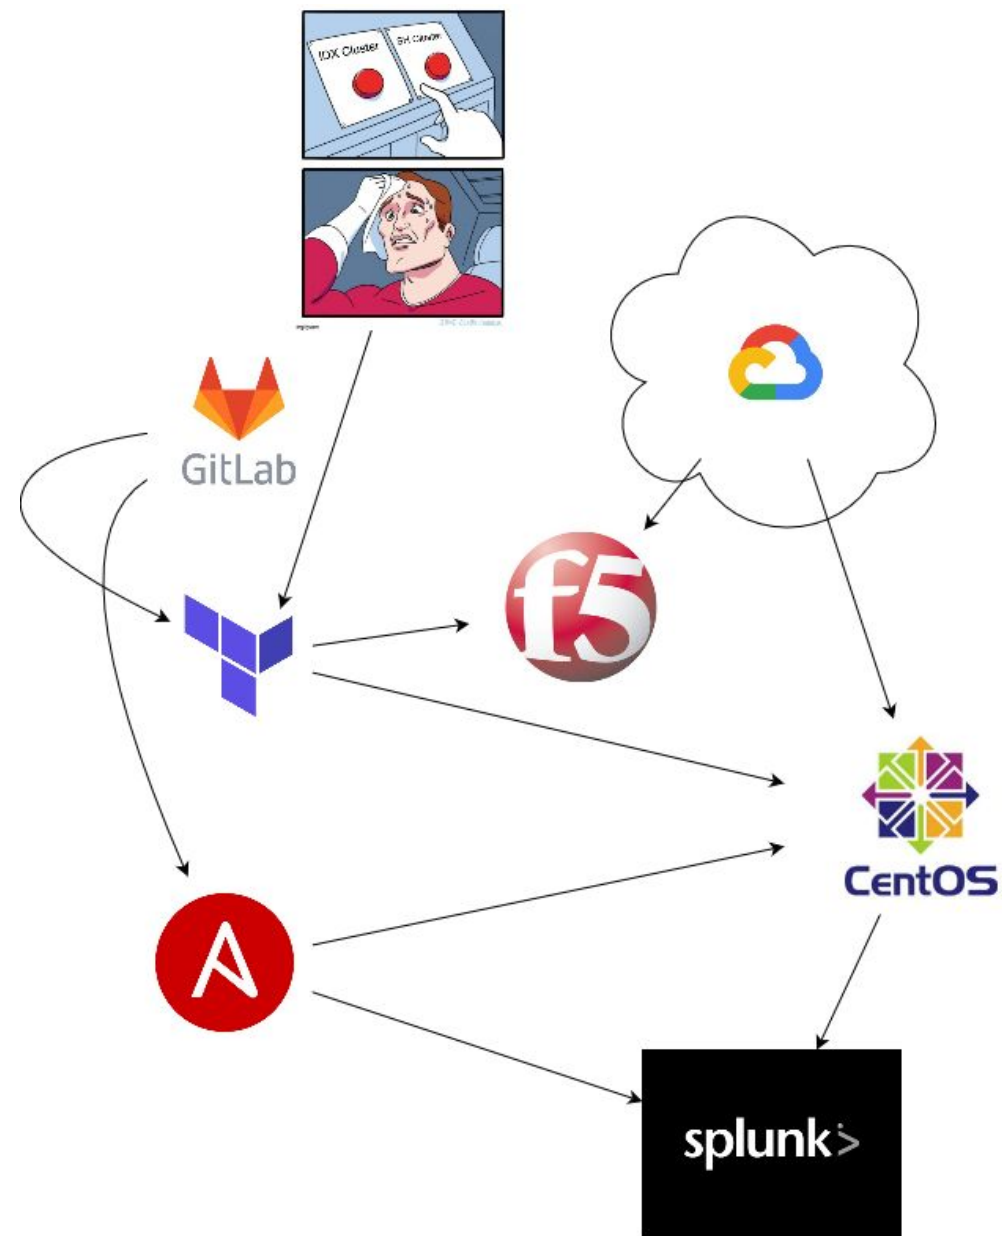


A little background:

- Build a Splunk environment in GCP
- Be lazy, automate as much as possible
- Create “single button” deployment.



The plan:



What is a "Derp Gun"

- In Online Gaming - A gun that causes a lot of damage with one shot *urbandictionary.com
- In this context a Derp Gun is a "one shot" Splunk deployment.



Provisioning with a Derp Gun in GCP

- Streamline Provisioning of New Resources in GCP
- Automate complicated setups such as Load Balancers
- Rapidly provision hosts, disk, network infrastructure



Configuration with a Derp Gun in GCP

- Automate configuration of Splunk hosts by role.
- Maintain configurations as repos in Git.
- Allow for new Splunk hosts to be configured as quickly as they can be provisioned.
- Maintain consistent configurations across the entire deployment.



The Tools of the Trade

- GCP Console
- Ansible AWX
- Terraform
- Gitlab
- Vault



- Obviously: Cloud Infrastructure
- IAM
- Autoscaling of Instance Groups



Google Cloud

- Create infrastructure as code
- Providers and Resources
- Available by default in GCP
- Works with Azure, AWS, AliCloud, etc
- Idempotent
- Can be used with anything that has an API: SignalFX, Dominos Pizza, etc



Terraform



- Perform Automated Tasks
- Configuration of Hosts
- Deployment of Apps, Indexes

Ansible



ANSIBLE




- Store configurations used by Terraform and Ansible
- Revision Control
- Collaboration
- CI/CD
- More Dev/Ops buzzwords..



Gitlab, Github, etc

Vault

- Stores secrets
- Works with Terraform, Ansible
- Token based access
- Integrates with Cloud KMS, IAM
- Can act as a Certificate Authority
- Can perform encryption services




Setting up a Terraform Environment

- Use bucket storage to allow tfstate to be shared
- Enable versioning on the bucket to retain backups of the state file in case of corruption
- Use a cloud ACL to limit access to the state file
- Bucket is accessible from the cloud shell
- Bucket provides state locking



Simple Terraform

Example: State Bucket



```
terraform {  
  backend "gcs" {  
    bucket = "ikea-splunk-prod-tf-state-bucket"  
    prefix = "terraform/tfstate"  
  }  
}
```

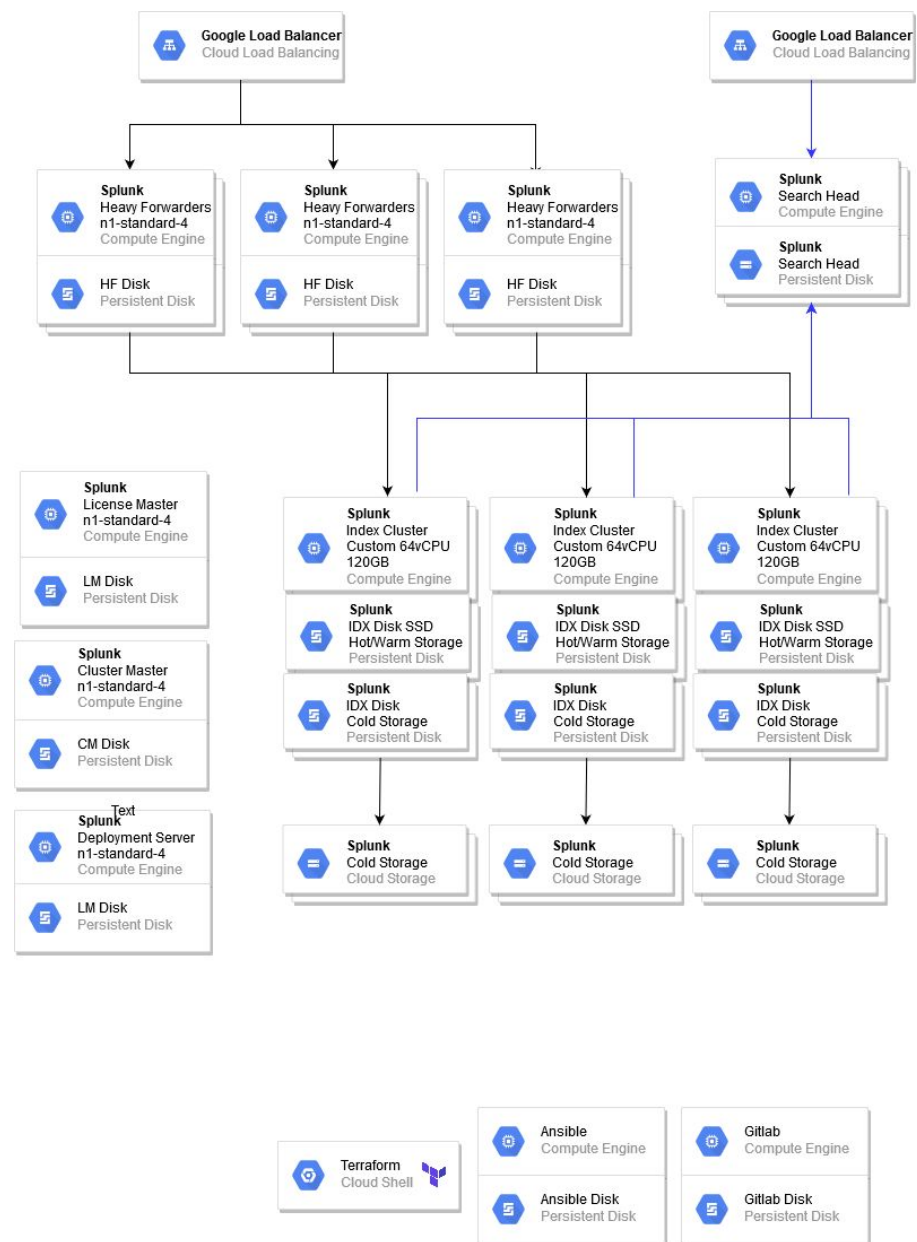
Start with a sturdy bucket.

- Bucket to store tfstate file
- Keep last 50 versions of the file
- ACL to protect access

```
// GCP Provider Config
```

```
provider "google" {  
  // Credentials not needed as this tf is run from a cloud shell with Terraform.  
  project = "${var.provider_project}"  
  region  = "${var.provider_region}"  
  zone    = "${var.splunk_gcp_zone}"  
}  
  
resource "google_storage_bucket" "ikea-splunk-gcp-tf-state" {  
  name      = "${format("%s", "${var.provider_project}-tf-state-bucket")}"  
  versioning {  
    enabled = "true"  
  }  
  location = "eu"  
  force_destroy = "true"  
  lifecycle_rule {  
    action {  
      type = "Delete"  
    }  
    condition {  
      NumberOfNewerVersions = "50"  
    }  
  }  
}  
  
resource "google_storage_bucket_acl" "image-store-acl" {  
  bucket = "${google_storage_bucket.ikea-splunk-gcp-tf-state.name}"  
  predefined_acl = "projectprivate"  
}
```


Provisioning the Needed Resources



Terraforming a Single Host

- Variables declared in tfvars file
- Tags are used by other resources such as load balancers, instance groups and firewalls
- Labels are used by Ansible
- Boot disk is created by default, destroyed if the host is changed.

```
resource "google_compute_instance" "indexer" {  
  count      = "${var.indexer_count}"  
  name       = "${format("%s", "idx${count.index}-${var.cluster_name}-gcp-${var.region}})"  
  machine_type = "${var.index_machine_type}"  
  tags       = [ "${var.cluster_name}", "index", "ssh" ]  
  labels = {  
    environment = "${var.cluster_name}"  
    function     = "idx"  
    ansible_group = "splunk_${var.cluster_name}_idx"  
    splunk_group = "splunk_${var.cluster_name}"  
  }  
  
  boot_disk {  
    initialize_params {  
      image = "${var.index_image}"  
      type  = "${var.index_disk_type}"  
    }  
  }  
}
```

Provisioning Storage

- Disks for Splunk, Warm and Cold
- Persistent Disks are not destroyed by host changes.
- Disks can be increased in size and filesystems resized with Ansible
- Prevent_destroy lifecycle function doesn't really work as expected.

```
// Index Splunk
```

```
resource "google_compute_disk" "idx-hot-warm-disk" {  
  count = "${var.indexer_count}"  
  name  = "${format("idx${count.index}-${var.cluster_name}-gcp-${var.region}-warm")}"  
  type  = "${var.index_hot_warm_disk_type}"  
  size  = "${var.index_hot_warm_disk_size}"  
  labels = {  
    environment = "splunk"  
  }  
  physical_block_size_bytes = 4096  
  lifecycle {  
    // prevent_destroy = true  
  }  
}
```

```
// Index Cold
```

```
resource "google_compute_disk" "idx-cold-disk" {  
  count = "${var.indexer_count}"  
  name  = "${format("idx${count.index}-${var.cluster_name}-gcp-${var.region}-cold")}"  
  type  = "${var.index_cold_disk_type}"  
  size  = "${var.index_cold_disk_size}"  
  labels = {  
    environment = "splunk"  
  }  
  physical_block_size_bytes = 4096  
  lifecycle {  
    // prevent_destroy = true  
  }  
}
```

Attaching Provisioned Disks

- Declared in compute resource.
- Name declared in configuration is also visible in /dev/disk/by-id/

```
boot_disk {  
  initialize_params {  
    image = "${var.index_image}"  
    type  = "${var.index_disk_type}"  
  }  
}  
  
attached_disk {  
  source      = "${format("idx${count.index}-${var.cluster_name}-gcp-${var.region}")}-splunk"  
  device_name = "${format("idx${count.index}-${var.cluster_name}-gcp-${var.region}")}-splunk"  
}  
  
attached_disk {  
  source      = "${format("idx${count.index}-${var.cluster_name}-gcp-${var.region}")}-warm"  
  device_name = "${format("idx${count.index}-${var.cluster_name}-gcp-${var.region}")}-warm"  
}  
  
attached_disk {  
  source      = "${format("idx${count.index}-${var.cluster_name}-gcp-${var.region}")}-cold"  
  device_name = "${format("idx${count.index}-${var.cluster_name}-gcp-${var.region}")}-cold"  
}
```


Compute Host Network Settings

- Public and Private subnets and interfaces
- Public IPs use NAT
- Firewall rules set globally upstream

```
network_interface {  
  subnetwork = "${var.pub_subnetwork}"  
  
  access_config {  
    nat_ip = "${google_compute_address.static[count.index].address}"  
  }  
}  
  
network_interface {  
  subnetwork = "${var.pri_subnetwork}"  
  access_config {  
  }  
}  
}
```

Using Modules to Create Clusters


- Variables declared in Module
- Module is a subfolder of the project
- Indexer_count builds number of hosts

```
// Index Clusters
```

```
module "alpha_index_cluster" {  
  source           = "./modules/index_cluster"  
  project          = "${var.provider_project}"  
  region           = "${var.provider_region}"  
  cluster_name     = "alpha"  
  indexer_count    = "3"  
  //index_machine_type = "custom-64-122880"  
  index_machine_type = "n1-standard-2"  
  index_image       = "gce-uefi-images/centos-7"  
  index_disk_type   = "pd-ssd"  
  index_disk_size   = "30"  
  index_hot_warm_disk_type = "pd-ssd"  
  index_hot_warm_disk_size = "50"  
  index_cold_disk_type = "pd-standard"  
  index_cold_disk_size = "50"  
  cm_machine_type   = "n1-standard-2"  
  cm_image          = "gce-uefi-images/centos-7"  
  cm_disk_type      = "pd-ssd"  
  cm_disk_size      = "30"  
  pub_subnetwork    = "${google_compute_subnetwork.public_subnet.name}"  
  pri_subnetwork    = "${google_compute_subnetwork.private_subnet.name}"  
  create_cm         = "1"  
  splunk_disk_type  = "pd-ssd"  
  splunk_disk_size  = "30"  
  dns_name_dmz      = "${google_dns_managed_zone.ikea-splunk-prod.dns_name}"  
  dns_name_int       = "${google_dns_managed_zone.ikea-splunk-prod-int.dns_name}"  
  managed_zone_dmz   = "${google_dns_managed_zone.ikea-splunk-prod.name}"  
  managed_zone_int   = "${google_dns_managed_zone.ikea-splunk-prod-int.name}"  
  num_reserved_ips   = 32  
}
```

Workflow with Vault

- Gitlab Configuration
- Vault Secrets
- Terraform Provisioning
- Ansible Configuration



Generate Secrets

- Generate Splunk Secrets
- Easily done in Terraform
- Don't ask about randomness

// Generate Vault Secrets

```
resource "random_password" "discover_pass4SymmKey" {  
  length = 64  
  special = true  
  override_special = "./"  
}  
  
resource "random_password" "pass4SymmKey" {  
  length = 64  
  special = true  
  override_special = "./"  
}  
  
resource "random_password" "admin_passwd" {  
  length = 64  
  special = true  
  override_special = "./"  
}
```


Store Secrets in Vault

- Write config data to Vault
- Used later by Ansible for configuration
- Minimize configuration required to build a new Splunk cluster.

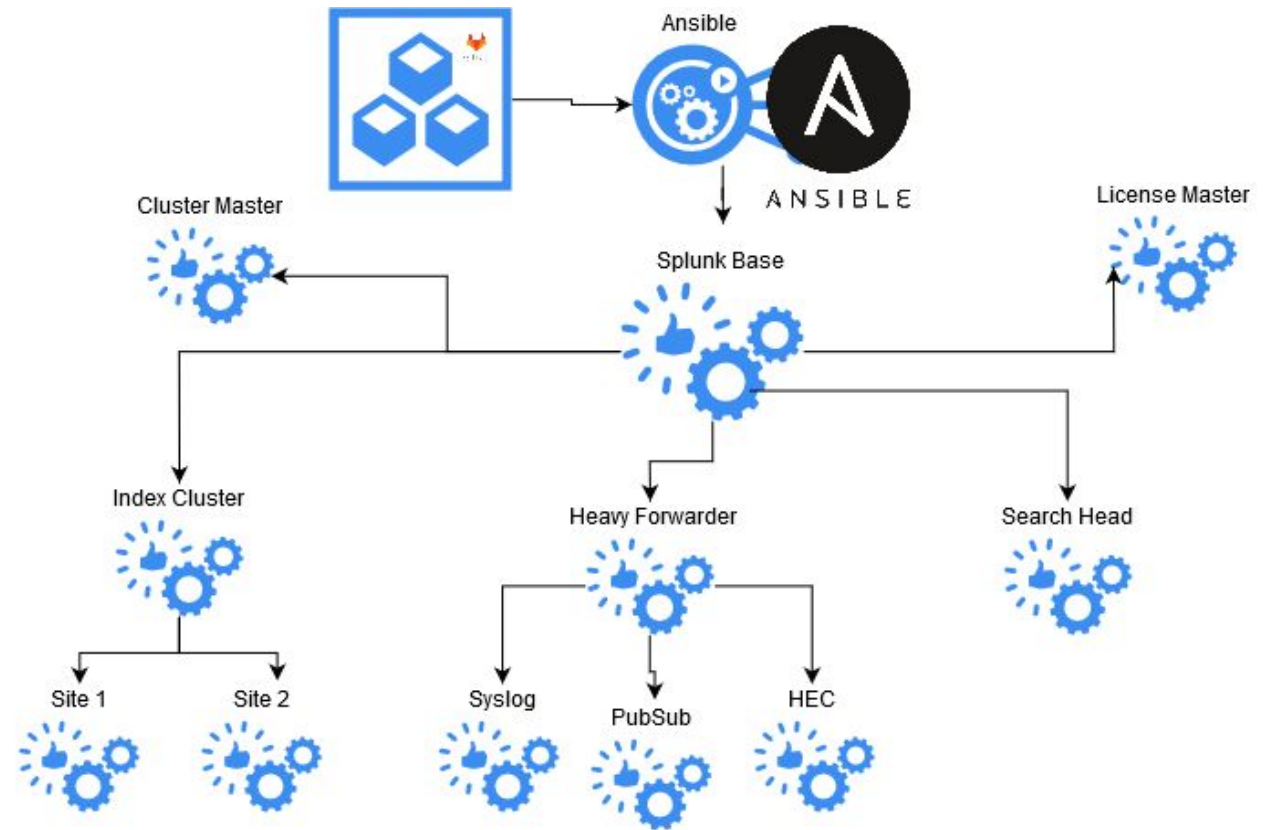
```
resource "vault_generic_secret" "index_cluster_vault" {
  path = "/secret/splunk/${var.cluster_name}"

  data_json = <<EOT
  {
    "splunk_cluster_discovery_pass4SymmKey": "${random_password.discover_pass4SymmKey.result}",
    "splunk_cluster_pass4SymmKey": "${random_password.pass4SymmKey.result}",
    "splunk_cluster_sslPassword": "${data.vault_generic_secret.sslPassword.data["sslPassword"]}",
    "splunk_current_admin_password": "changeme",
    "splunk_default_admin_password": "${random_password.admin_passwd.result}",
    "splunk_ds": "${var.ds_name}",
    "splunk_hf_label": "${var.cluster_name}",
    "splunk_ixc_label": "${var.cluster_name}",
    "splunk_lm": "${var.lm_name}",
    "splunk_master_ip": "${google_dns_record_set.cm-int.0.name}",
    "splunk_old_admin_password": "changeme",
    "splunk_serverCertName": "splunkmanagementdata.v4.pem",
    "splunk_sslRootCAPathName": "IKEA-CA-bundle.v4.pem"
  }
  EOT
}
```

GREAT PITCH, AMAZING PRODUCT

I'M OUT

Ansible Workflow



Ansible Playbooks

- Uses Labels in GCP set by Terraform
- Applies to Roles in Ansible



```
1 ---
2 # Configure all splunk nodes with common configuration
3 - name: Setup base config on Splunk servers
4   hosts: "{{ servers | default('splunk_alpha') }}"
5   roles:
6     - { role: common, tags: ['common', 'common-base'] }
7     - { role: splunk-base, tags: ['splunk-base', 'common-base'] }
8
9 # Configure the Cluster Master
10 - name: Setup config on Splunk Cluster Master
11   hosts: "{{ cm_servers | default('splunk_alpha_cm') }}"
12   roles:
13     - { role: splunk-cluster-master, tags: splunk-cluster-master }
14     - { role: splunk-web-interface, tags: splunk-web-interface }
15
16 # Add the Peer Nodes to the cluster
17 - name: Setup config on Splunk Peer Nodes
18   hosts: "{{ idx_servers | default('splunk_alpha_idx') }}"
19   roles:
20     - { role: splunk-peer-nodes, tags: splunk-peer-nodes }
21
22 # Install Heavy Forwarder for Adapters
23 - name: Heavy Forwarder for Adapters
24   hosts: "{{ hf_adapter_servers | default('splunk_alpha_hf_adapter') }}"
25   roles:
26     - { role: splunk-heavy-forwarder, tags: splunk-heavy-forwarder }
27     - { role: splunk-heavy-forwarder-adapter, tags: splunk-heavy-forwarder-adapter }
28     - { role: splunk-web-interface, tags: splunk-web-interface }
29
30 # Install Heavy Forwarder for UF traffic
31 - name: Heavy Forwarder for traffic to 9997 config
32   hosts: "{{ hf_9997_servers | default('splunk_alpha_hf_in') }}"
33   roles:
34     - { role: splunk-heavy-forwarder, tags: splunk-heavy-forwarder }
35     - { role: splunk-heavy-forwarder-9997, tags: splunk-heavy-forwarder-9997 }
36
37 # Install Heavy Forwarder HEC
38 - name: Heavy Forwarder HEC config
39   hosts: "{{ hf_hec_servers | default('splunk_alpha_hf_hec') }}"
40   roles:
41     - { role: splunk-heavy-forwarder, tags: splunk-heavy-forwarder }
42     - { role: splunk-heavy-forwarder-hec, tags: splunk-heavy-forwarder-hec }
```



Ansible Roles



Name
..
cluster-uninstall/tasks
common/tasks
git
splunk-base
splunk-cluster-master
splunk-deployment-server
splunk-heavy-forwarder
splunk-heavy-forwarder-9997
splunk-heavy-forwarder-adapter
splunk-heavy-forwarder-hec
splunk-license-master
splunk-peer-nodes
splunk-web-interface



Variables and Secrets

```
1  ---
2  # Override default values
3  # Specific configuration for global staging
4  deploy_stage: "alpha"
5
6  # Packages
7  # Use wget to grab the latest package from splunk.com
8  # Place file in roles folder under -> splunk-base/files
9  splunk_package_file: 'splunk-7.2.4-8a94541dcfac.x86_64.rpm'
10 splunk_version: '7.2.4'
11 splunk_base: '/opt/splunk'
12 splunk_vg_name: vg_splunk
13 splunk_lv_name: lv_splunk
14 file_system: xfs
15
16 # Set mountpoints for buckets
17 idx_warm_path: "{{ splunk_base }}/var/lib/splunk"
18 idx_cold_path: "{{ splunk_base }}/var/lib/splunk/cold"
19 warm_vg_name: vg_warm
20 warm_lv_name: lv_warm
21 cold_vg_name: vg_cold
22 cold_lv_name: lv_cold
23
24 # Use multisite or single site cluster config
25 #splunk_cluster_pass4SymmKey: 'located in secrets file'
26 use_multisite_config: "false" #set to false when not using
27
28 # Multisite config options
29 site: "site0"
30 cluster_label: "alpha_Cluster"
31 search_site: "site0"
32 available_sites: "site0"
```

```
29 alpha:
30     splunk_ixc_label: "alpha"
31     splunk_hf_label: "alpha"
32     splunk_master_ip: "cm-alpha-gcp-europe-west1.ikea-splunk-prod.int"
33     splunk_ds: 'ds0-gcp-europe-west1.ikea-splunk-prod.int'
34     splunk_lm: 'lm.ikea-splunk-prod.int'
35     splunk_default_admin_password: 'changeme'
36     splunk_old_admin_password: 'changeme'
37     splunk_current_admin_password: 'changeme'
38     splunk_cluster_pass4SymmKey: 'someRandomKey'
39     splunk_cluster_discovery_pass4SymmKey: 'differentRandomKey'
40     ssl:
41         splunk_cluster_sslPassword: 'IDon'tThinkSo'
42         splunk_sslRootCAPathName: "IDon'tThinkSo"
43         splunk_serverCertName: "IDon'tThinkSo"
44
```



Push and Pull to Git

- Use Ansible to create projects in GitLab from templates
- Pull config from project if it already exists
- One place for all the current configs



```
15 - name: Create Gitlab Project
16   gitlab_project:
17     server_url: "{{ git_server }}"
18     api_token: "{{ git_token }}"
19     group: "{{ git_group }}"
20     name: "{{ customer_name }}_{{ stage.splunk_ixc_label }}_idx"
21     state: present
22     validate_certs: no
23   delegate_to: localhost
24   run_once: true
25   register: gitCreate
26
27 - name: Populate new Git project
28   import_tasks: create_conf.yml
29   when: gitCreate.changed
30
31 - name: "Git: is it up-to-date?"
32   git:
33     repo: "{{ git_address }}:{{ git_group }}/{{ customer_name }}_{{ stage.splunk_ixc_label }}_idx"
34     dest: "{{ git_tmp }}/{{ customer_name }}_{{ stage.splunk_ixc_label }}_idx"
35     update: yes
36     version: master
37   register: git
38   delegate_to: localhost
39   run_once: true
40 - debug:
41   var: git
42
43 - name: "Create {{ customer_name }}_{{ stage.splunk_ixc_label }}_master_base/local directories"
44   file:
45     path: "{{ item }}"
46     owner: splunk
47     group: splunk
48     mode: u=rwX,g=rX,o-rwx
49     recurse: yes
50     state: directory
51   with_items:
52     - "{{ splunk_base }}/etc/apps/{{ customer_name }}_{{ stage.splunk_ixc_label }}_master_base,
53     - "{{ splunk_base }}/etc/apps/{{ customer_name }}_{{ stage.splunk_ixc_label }}_master_base,
```



Git Project


- Separate project for each cluster and function
- All Splunk config is in apps, nothing in system/local
- Using Splunk base apps standard



**ikea_alpha_hf** 
Project ID: 64

[Add license](#) [3 Commits](#) [1 Branch](#) [0 Tags](#) [246 KB Files](#)

master ikea_alpha_hf / +

 **Create templates**
Robert Johansson authored 1 minute ago

[Auto DevOps enabled](#) [Add README](#) [Add CHANGELOG](#)

Name
ikea_alpha_hf_9997_inputs/local
ikea_alpha_hf_all_inputs/default
ikea_alpha_hf_all_server/default
ikea_alpha_hf_deploymentclient/default
ikea_alpha_hf_hec_inputs/local
ikea_alpha_hf_outputs/default







What could possibly go wrong?



- Easy come, easy go
- Corrupt Terraform Statefile
- Terraform Lifecycle doesn't work if code is removed.
- Restarting the cluster all at once







Some assembly
required.



Bigger, Better, More Derpy

- Autoscaling
- Further Vault Integration
- Single Configuration Source
- Log All The Things to Splunk
- SmartStore for Cold Buckets
- Automagic PubSub
- Consul for Connectivity





memegenerator.net