



HANA EXPERT SERIES

HANA Heroes Expert Session

Title: Automating SAP Deployments
with Terraform and Ansible

Presenter: Darpan Patel

Proudly presented by **SAP HANA Heroes**

Kick Off Automated **SAP S/4HANA** Build



Project Thunder

- Join [here](#) if you are not a member so you don't miss updates regarding this topic and many more!
- Blog Post on this topic will be released beginning of Q2

SAP Jam This Group Search this group...

Home Groups Knowledge Base

Project Thunder: Overview AWS Azure GCP SAP BTP News and Events Relevant Services Tools Who's who

Project Thunder Hyperscaler Readiness

Project Thunder is a S/4HANA Move initiative to ensure IDG readiness for S/4HANA on Hyperscalers. This Jam site is a collection of resources that will be continually updated focused on all relevant topics like Team Enablement through persona-specific Learning Maps, Services Catalog, and Project Run-Time tools.

"Many customers I have spoken with have expressed the need for agility and quick time to value, made even more urgent by the COVID-19 pandemic. To achieve both, technology becomes even more important, and SAP can help meet those challenges. We've added significant enhancements to every major component of our Business Technology Platform to help customers across all industries overcome obstacles stemming from ubiquity of data, complexity of IT and business volatility." - Jürgen Müller, SAP CTO and Executive Board member

News & Updates

- CLF-01 – AWS Certified Cloud Practitioner Study Circle. Schedule, presentations and recordings are available [here](#)
- Our SAP Business Technology Platform (BTP) is the technical foundation of Intelligent Enterprise. More info in the [SAP BTP tab](#)
- AZ-120 – Planning and Administering Microsoft Azure for SAP Workloads. Presentations and recordings are available [here](#)
- Check out the [step-by-step instructions](#) to take advantage of free MS Azure Certification
- AZ-900 – Azure Fundamentals Study Circle. Presentations and recordings are available [here](#)
- Project Thunder is launching several one-pagers that shows project details and contact information. See [News and Event](#) for more information
- Learn more about SAP S/4HANA Deployments on hyperscaler platforms in [this blog post](#)
- Hyperscaler Support & information Resources can be found [here](#) (Note: 2951356)

Agenda

- Why learn these tools?
- Getting Started with Terraform
- Getting Started with Ansible
- Merging Terraform and Ansible



Why Terraform and Ansible?



SAP Implementations

- Implementation Issues:
 - Technical Implementations can become very lengthy
 - They are a bottleneck for the project
 - Involve a very manual process
 - Both Provisioning servers and applications
 - Outdated in today's quick paced and automated world
- How do we make implementations more efficient?
 - Automate
 - Provisioning Servers
 - Resources that attach onto those servers
 - The networks they reside in
 - OS configurations
 - Application Deployments





HashiCorp

Terraform



RED HAT[®]
ANSIBLE[®]
Automation

- Infrastructure as Code
 - Use to Deploy the Hosts for SAP application servers
 - Database server hosts
 - Storage
 - Networks, VPC's, Subnets, routers
 - Security Groups, NACL's
- Deploy applications
 - Install Packages for OS
 - Configure your file system
 - Mount Network locations
 - Install SAP HANA
 - Any task SWPM can do
 - System Copy, System Rename, Distributed System, Standard System

Getting Started with Terraform



Terraform Implementation Considerations

- Terraform is logically split into two main parts:
 - Terraform Core
 - Terraform Plugins
- Terraform is Agentless
- Terraform can be installed on any platform
- Installation Options
 - [Manual](#)
 - [OS Package Manager](#)
- Terraform Core and Plugins are written in the GO Language
- Terraform uses CRUD (create, read, update, and delete) API's to communicate with providers

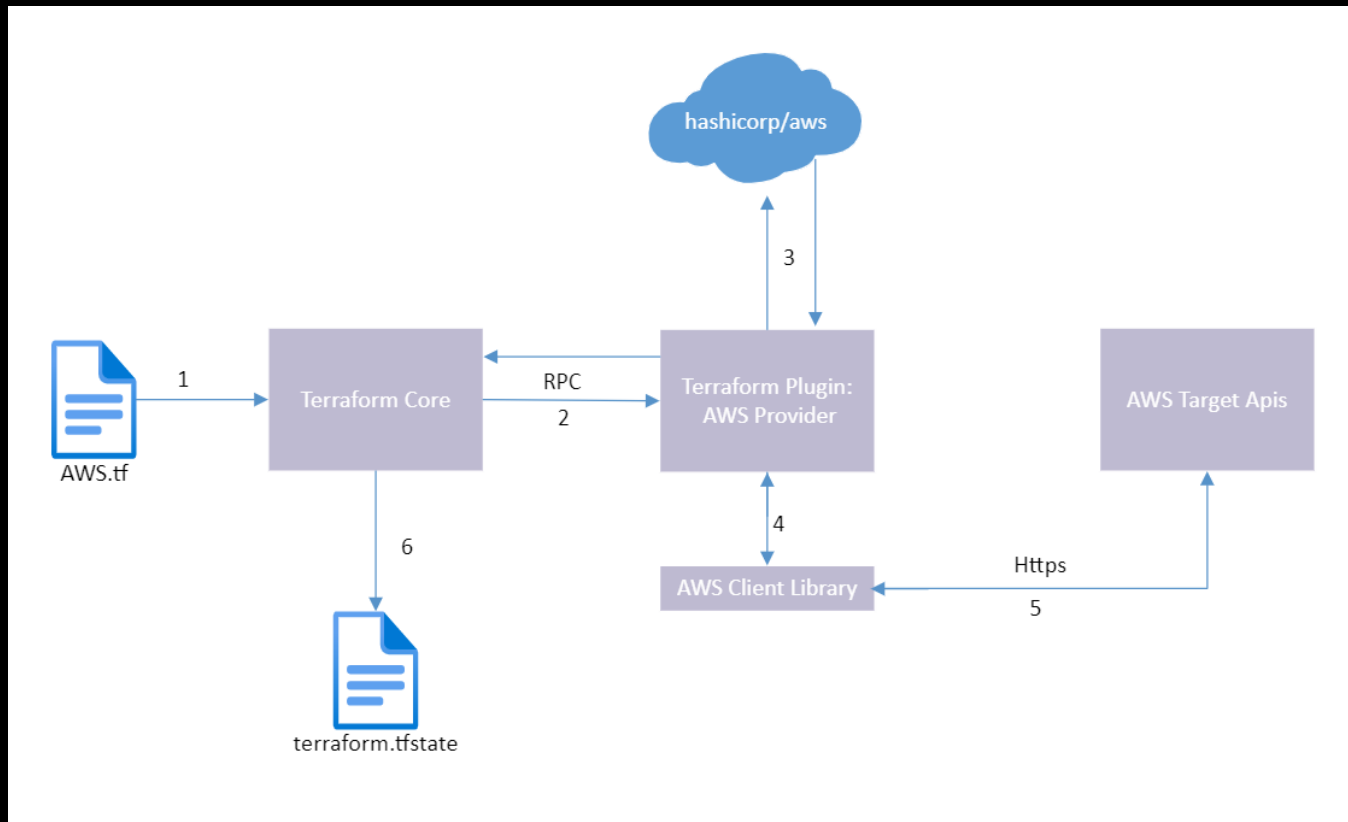


- Statically-compiled binary written in GO
- Terraform core is the command line tool
- Entry point for anyone using terraform
- Resource State Management
- Communicates with plugins



- Terraform plugins are executable binaries written in GO
- Plugins contain Providers and Provisioners
- Providers provide a service such as AWS

Terraform Execution Flow



- Executing a Terraform Script to provision in AWS:
 - 1 – Execute terraform against `AWS.tf`
 - 2 – Core makes Remote Procedure Call to Plugin to download AWS Provider and provide resources to provision
 - 3 – Plugin downloads AWS Libraries- ADD where the request goes
 - 4 – AWS Client library translates requests to API requests
 - 5 – API requests sent to AWS
 - 6 – Once provisioning is complete, state data is logged in `terraform.tfstate`

Terraform Providers

AWS.tf Terraform File:

```
1
2 terraform {
3   required_providers {
4     aws = {
5       source = "hashicorp/aws"
6       version = "3.31.0"
7     }
8   }
9 }
10
11 provider "aws" {
12   # Configuration options
13 }
```

- Terraform Providers:
 - To install a provider you must include provider configuration code into your terraform files
 - A list of available providers can be found [here](#) on the terraform website
 - Each Provider comes with its own set of documentation describing its resource types and their arguments
 - Terraform currently has 70 providers in their registry
 - You can create your own provider

Terraform Resources

Resource Template:

```
1 resource "<PROVIDER>_<TYPE>" "<NAME>" {
2     [CONFIG ...]
3 }
```

AWS.tf – Create an EC2 Instance:

```
15
16 resource "aws_instance" "example" {
17     ami = "ami-0c09927662c939f41"
18     instance_type = "t2.micro"
19     tags { name = "TESTVM" }
20 }
```

- Resources:
 - Resources are the most important element in the Terraform language
 - They represent some type of infrastructure object
 - virtual networks, compute instances, compute devices or DNS records.
 - Terraform is platform dependent
 - Each provider has different requirements for parameters
 - AWS requires – ami ID, and Instance_type

Terraform Variables

AWS.tf Terraform File with Variables :

```
15
16 resource "aws_instance" "example" {
17     ami = var.amiID
18     instance_type = var.instanceType
19     tags { name = var.tags}
20 }
21
22
```

Variables.tf File :

```
1 variable "amiID"{
2     type = string
3     default = "ami-0c09927662c939f41"
4 }
5
6 variable "instanceType"{
7     type = string
8     default = "t2.micro"
9 }
10
11 variable "Tags"{
12     type = string
13     default = "TEST_VM"
14 }
```

- Variables:
 - Variables can be used to make your terraform scripts re-usable
 - Avoid having to hardcode in your parameters by using variables
 - Easier to Maintain – Instead of maintaining script file, you just need to maintain variables file
- Create Variables File:
 - Create a file in the same directory as your provisioning script and title it variables

Terraform Best Practices for File Structure

Example File Structure :


```
darpan-test:/terraform # tree -l
.
├── dev
│   ├── AWS.tf
│   └── variables.tf
├── prod
│   ├── AWS.tf
│   └── variables.tf
└── qa
    ├── AWS.tf
    └── variables.tf

3 directories, 6 files
darpan-test:/terraform # |
```

- File Structure:
 - Separate out clusters of Servers
 - An example would be breaking out your environments to include DEV, QA, and PROD
 - Each environment has its own execution script – AWS.tf
 - Can also create groups such as Active Directory Servers, DNS servers, Printer Servers, etc..
- Drawbacks to this:
 - Every time you upgrade or modify one AWS.tf, you will need to modify the others

Terraform Modules

Example File Structure with Modules :



- Modules
 - Don't need to update provisioning scripts in multiple locations
 - All resource provisioning scripts will be put into the module
 - Useful when managing multiple environments with different types of servers
 - Modules can be shared to the community [here](#)
- Use Case
 - You want to build out another application server for your S/4 Dev Environment
 - Update `/dev/AWS.tf` to include the module `/modules/s4-server/application`
 - To execute – cd into `/dev` and run: *terraform apply*

Getting Started with Ansible



Ansible Implementation Considerations

- Types of Resources in Ansible
 - Control Node
 - Managed Node
- Ansible is Agentless
- Installation Options
 - Install using OS package Manager
 - [RHEL](#)
 - [SUSE](#)
 - Install with Python Package Manager [pip](#)
- Control Node uses ssh to communicate with Managed Nodes
- Ansible Modules/Libraries are written in python

Control Node



- Cannot be a Windows System
- Can be a laptop, shared desktop, or server
- RHEL, SUSE, macOS, Ubuntu, other linux operating systems
- Control Node should be located on the same network as Managed Nodes
- Must have python installed, python 3 is recommended

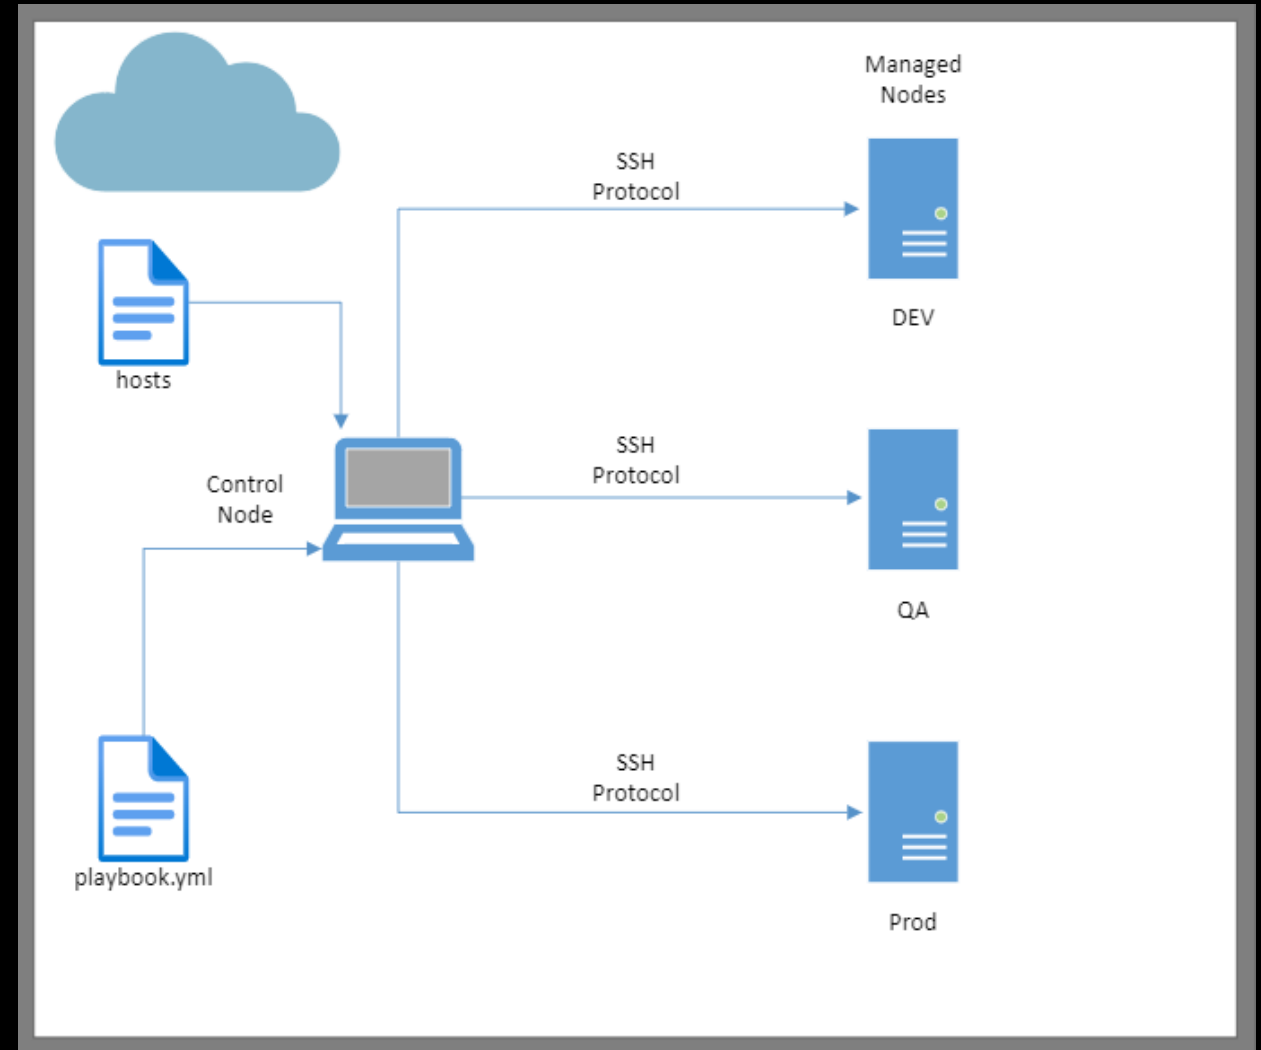
Managed Node



- Can be any operating system, including Windows
- Must have python installed, python 3 is recommended

Ansible Execution Flow

- Authentication
 - SSH Keys or Passwords
- Hosts File
 - Contains the list of Managed Nodes
- Playbook
 - Contain Tasks to execute against the hosts in Hosts file



Ansible Hosts File

- Hosts File
 - Host file can be written in YAML or INI format
 - Managed nodes can use FQDN, Short Name, or IP Addresses
 - Group out your Servers for segregation of environments
 - Make groups of groups using :children suffix to make your execution calls more robust
 - Execute Ansible Playbooks against a server, a group of servers, and groups of groups of servers
 - More Information on Hosts files [here](#)

```
[sapservers:children]
1 [sapservers:children]
2 dev
3 qa
4 prod
5
6 #####
7
8 [dev:children]
9 app_dev
10 db_dev
11
12 [app_dev]
13 sapapp1
14 sapapp2
15
16 [db_dev]
17 hanadb
18
19 #####
20
21 [qa:children]
22 app_qa
23 db_qa
24
25 [app_qa]
26 10.xx.xxx.xx
27
28 [db_qa]
29 10.xx.xxx.xx
```

Ansible Playbook Files

- Playbook Files
 - Playbooks are written in YAML format
 - You must specify a host or a group of hosts to run the play against
 - Define your task by providing a name, module and parameters for each task
 - Playbooks define the execution flow of your Ansible tasks
- Execute Playbook command:
 - `Ansible-playbook <playbook_name>`
- Playbook Drawbacks
 - Can get messy when you have too many tasks

```
1 - hosts: dev
2   tasks:
3     - name: Install SAP Hana Prerequisite Packages for SUSE Linux 15
4       zypper:
5         name: libgcc_s1
6         state: present
7
8     - name: Install SAP Hana Prerequisite Packages for SUSE Linux 15
9       zypper:
10        name: libstdc++6
11        state: present
12
13    - name: Install SAP Hana Prerequisite Packages for SUSE Linux 15
14      zypper:
15        name: libatomic1
16        state: present
17
18
```

Ansible Roles

- Roles
 - Provide a clean way to organize your Ansible scripts and not overpopulating your playbooks
 - Treat roles like functions in programming
 - Should perform one task or a group of similar tasks
 - Roles are reusable and can be shared to your team or the [Ansible Galaxy Community](#)
- Role Creation
 - In the same directory your playbook can be found create a folder called roles and change directory into it
 - Run - `ansible-galaxy init <Role_Name>`
 - Folder structure for role is created

New look Playbook:

```
1 - hosts: dev
2   roles:
3     - { role: sap-configure }
4
5
```

Role task file:

```
1 ---
2 # tasks file for sap-configure
3
4 - name: Install SAP Hana Prerequisite Packages for SUSE Linux 15
5   zypper:
6     name: libgcc_s1
7     state: present
8
9 - name: Install SAP Hana Prerequisite Packages for SUSE Linux 15
10  zypper:
11    name: libstdc++6
12    state: present
13
14 - name: Install SAP Hana Prerequisite Packages for SUSE Linux 15
15  zypper:
16    name: libatomic1
17    state: present
18
```

Ansible Variable Files

- Variable File
 - Use variables to manage differences between systems
 - Prevents you from having to hard code in your playbook and role files
 - Create variables with YAML syntax, lists, or dictionaries
 - Define variables in playbooks, hosts file, roles, at the command line, or in separate vars files
 - Make all variables in vars file useable in your playbooks and roles
- Variable Drawbacks
 - Sensitive variables need to be secured as they are visible to users

Playbook File calling vars file:

```
1 - hosts: dev
2   vars_files:
3     - /etc/ansible/vars/ansible_vars.yml
4   vars:
5     ansible_ssh_private_key_file: "{{ private_key }}"
6   roles:
7     - { role: sap-configure }
8
9
```

Vars File:

```
1 #Secret Variables for Playbook
2 {
3   "private_key": "/etc/ansible/dpid_rsa"
4 }
5
```

Ansible Vault

- Vault
 - Encrypt and Decrypt your sensitive files with Ansible Vault
 - Before files are encrypted you must provide a password
 - When executing playbooks, if an ansible vault encrypted file exists, you will be prompted for the vault password
 - Files are decrypted during runtime only for Ansible to use in execution
 - Create multiple vault passwords for encrypting different sets of files
- Vault Encrypt Command
 - `ansible-vault encrypt <file_path1> <file_path2>`

Vars file before encryption:

```
1 #Secret Variables for Playbook
2 {
3     "private_key": "/etc/ansible/dpid_rsa"
4 }
5
```

Vars file after encryption:

```
1 $ANSIBLE_VAULT;1.1;AES256
2 31333532396633343931643130613139616666663636333339643465663864346336363656
3 3936313139636563663361643763346335323634613832320a3833616262613335356361323
4 326364353639303961353032363466616631636565363763373034353861373930613238623
5 3535363134366261640a6234306639376362613433363764386232666337353636633831663
6 643631333939306363316336386366306162643234646361623261613434666236363136346
7 386230623335663632366365626531613539366339303336613431656663373838303262643
8 616439653333373262663265333837666332356366393439373162626561313335633433396
9 393265626437616632626636306163613832656131343038376664383339623230333231343
10 613432336137623234373230346566393263636239303161306534626237336564366539383
11 333535656537313537343037393064376664343233623963326538373532306135373366353
12 643231633064316331313530386437356561656136333336306536373637333137323862326
13 316430643238643566383263373939353262333165613632353133326438633462316461313
14 613132666265666466323336626239663430623864636337626332356363626439346130333
15 663532333435306439663331323838356562383963383833663165353432393138623237383
16 383836646164316139306237636631313263353330373135633263646634386363353531316
17 39306434653635656366463336439313534356333366433626530383533386638613966366
18 34643638383830663433396661636534623665306362653435393936653537643930
19
```


Ansible File Structure

- File Structure
 - Ansible can be run from anywhere
 - Bare minimum you need a playbook and hosts file
 - Group Vars gives you the ability to associate variables with a specific group of hosts
 - Host Vars gives you the ability to associate variables to specific hosts

Example of Ansible File Structure:


```
darpan-test:/ansible # tree -l
.
├── group_vars
│   ├── dev.yml
│   ├── prod.yml
│   ├── qa.yml
│   └── sapservers.yml
├── host_vars
│   ├── 10.47.106.205.yml
│   ├── hanadb.yml
│   └── sapapp1.yml
├── hosts
├── roles
│   └── sap-configure
│       └── tasks
│           └── main.yml
├── sap-configure.yml
├── vars
│   └── ansible_vars.yml
└──
```

6 directories, 11 files
darpan-test:/ansible # pwd
/ansible



Use Ansible to Call Terraform

Combining Terraform and Ansible File Structures:



sap-configure.yml calls Terraform:

```
1 #####
2 - hosts: localhost
3   tasks:
4     - name: init terraform
5       shell: terraform init
6       args:
7         chdir: "/etc/ansible/terraform/dev"
8
9     - name: apply terraform script
10      terraform:
11        project_path: "/etc/ansible/terraform/dev"
12        state: present
13 #####
14 - hosts: dev
15   vars_files:
16     - /ansible/vars/ansible_vars.yml
17   vars:
18     ansible_ssh_private_key_file: "{{ private_key }}"
19   roles:
20     - { role: sap-configure }
```

How Ansible Performs SAP Installations

Hdblcm In Batch Mode and SWPM

Unattended:

- Config File needs to be maintained with parameters used for HANA and S/4 Installations
- You can write a playbook that picks up this file and inputs it into an hdblcm command or sapinst command
- No need to reinvent the wheel – download the following ansible roles
 - [SAP HANA](#)
 - [SAP S/4HANA](#)
- For SWPM, follow SAP Note - [2230669](#) to generate config files

Config File for SWPM:

```
# Standard system with AS ABAP only: ASCS instance number. Leave empty for default.
NW_CI_Instance.ascsInstanceNumber = {{ sap_s4hana_deployment_ascs_instance_nr }}

# Standard system with AS ABAP only: Virtual host name for the ASCS instance. Leave empty for default.
NW_CI_Instance.ascsVirtualHostname = {{ ansible_hostname }}

# Instance number of the primary application server instance. Leave empty for default.
NW_CI_Instance.ciInstanceNumber = {{ sap_s4hana_deployment_pas_instance_nr }}
```