Chapter 1: The Impact of Cloud On Networking

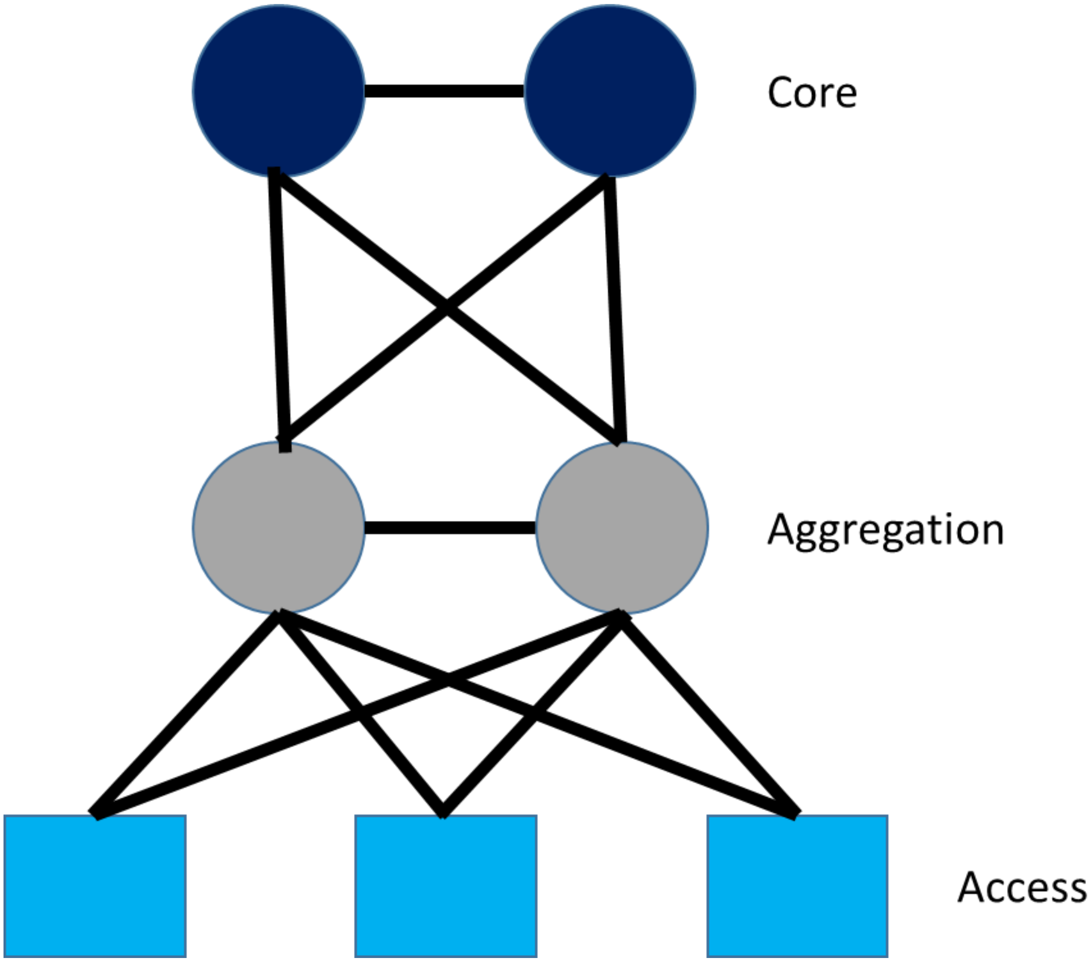

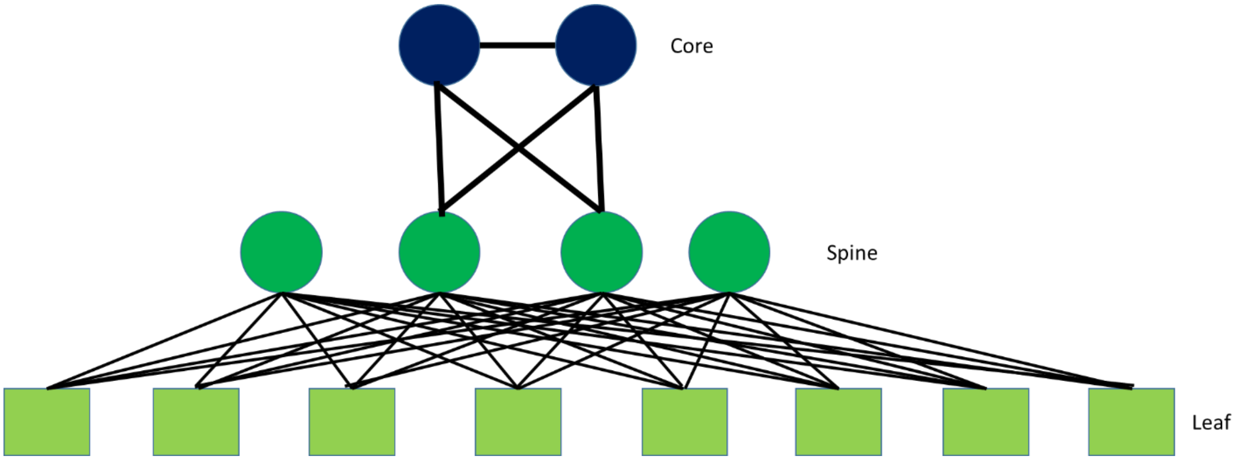

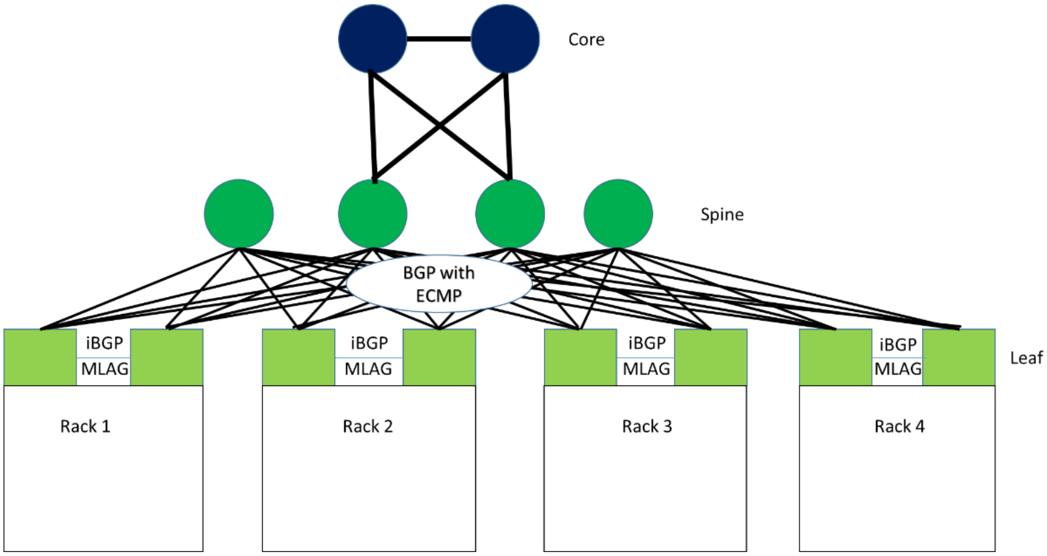

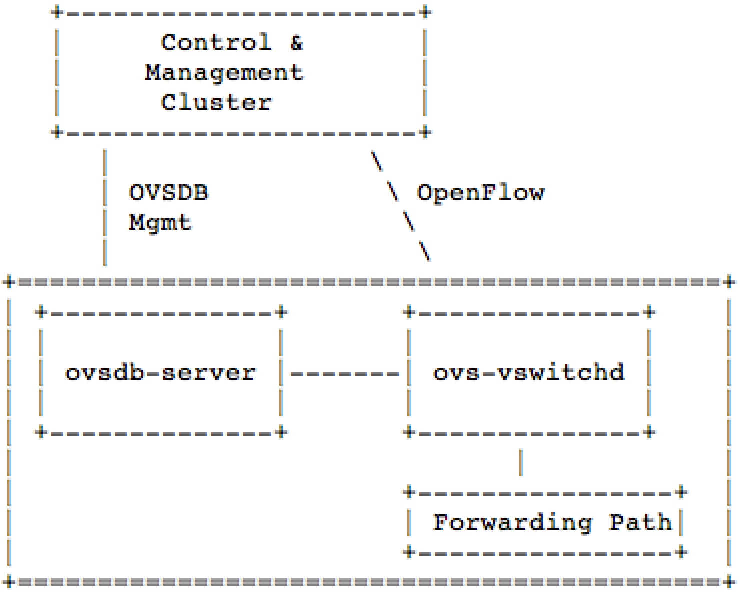

```
        +-----------------------+
        |      Control &        |
        |     Management        |
        |      Cluster          |
        +-----------------------+
           |                \
           | OVSDB           \ OpenFlow
           | Mgmt             \
           |                   \
+==============================================+
|  +--------------+      +--------------+      |
|  |              |      |              |      |
|  | ovsdb-server |------| ovs-vswitchd |      |
|  |              |      |              |      |
|  +--------------+      +--------------+      |
|                               |              |
|                        +----------------+    |
|                        | Forwarding Path|    |
|                        +----------------+    |
+==============================================+
```

AT&T Canonical Hewlett Packard Enterprise IBM Intel Rackspace

Red Hat, Inc. SUSE

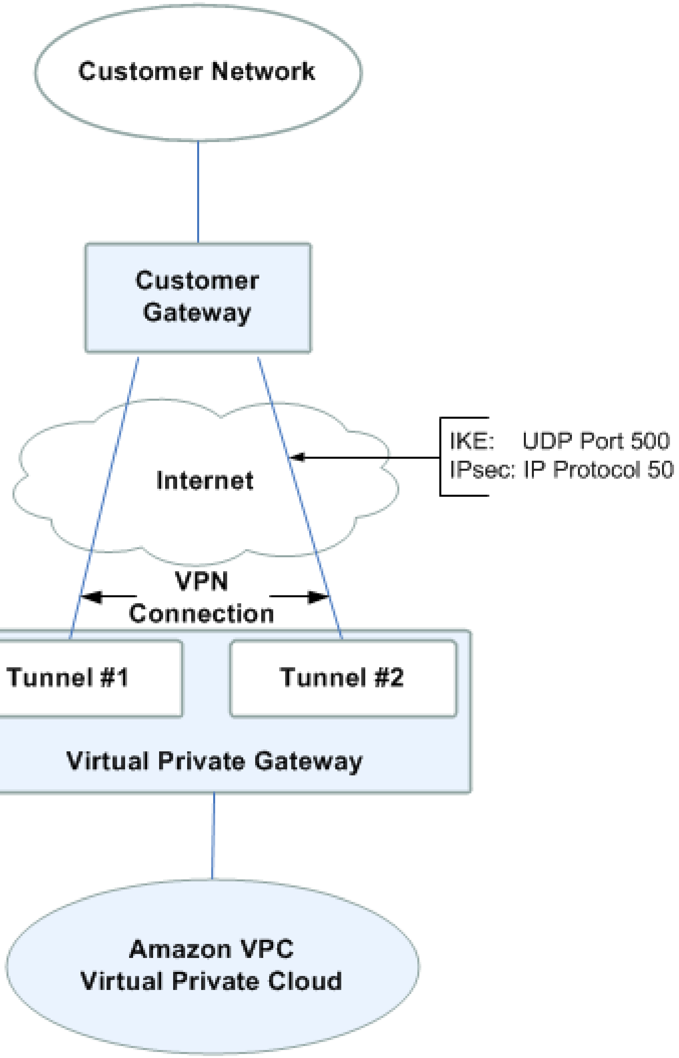
Customer Network
Customer Gateway
Internet
IKE: UDP Port 500
IPsec: IP Protocol 50
VPN Connection
Tunnel #1
Tunnel #2
Virtual Private Gateway
Amazon VPC
Virtual Private Cloud

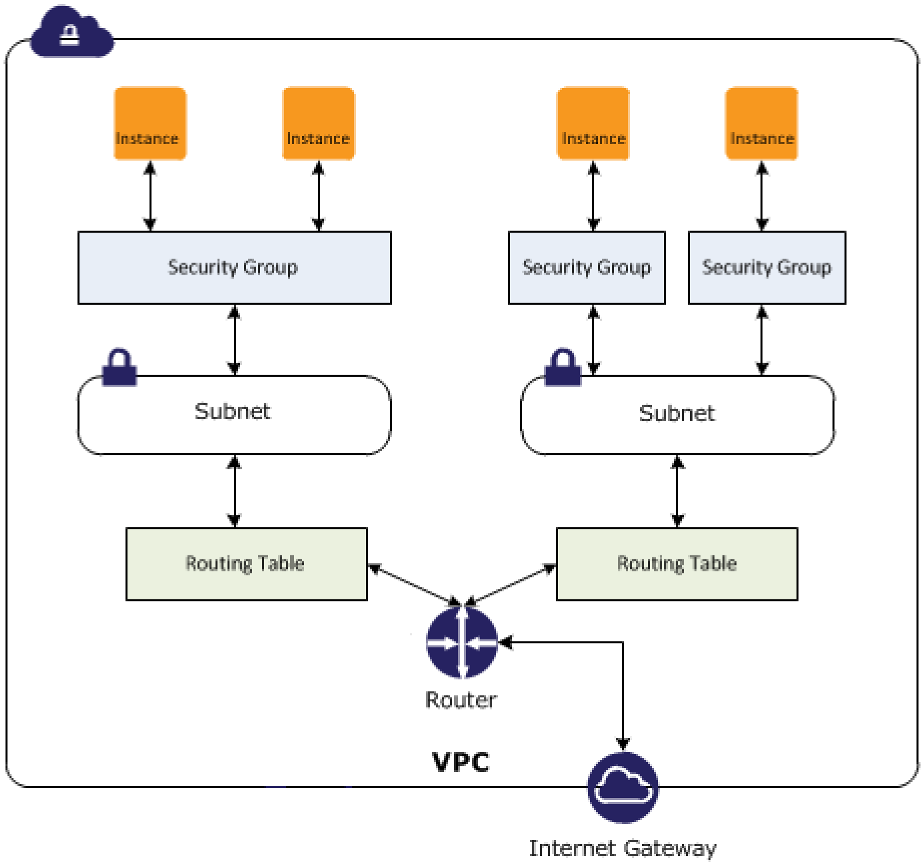

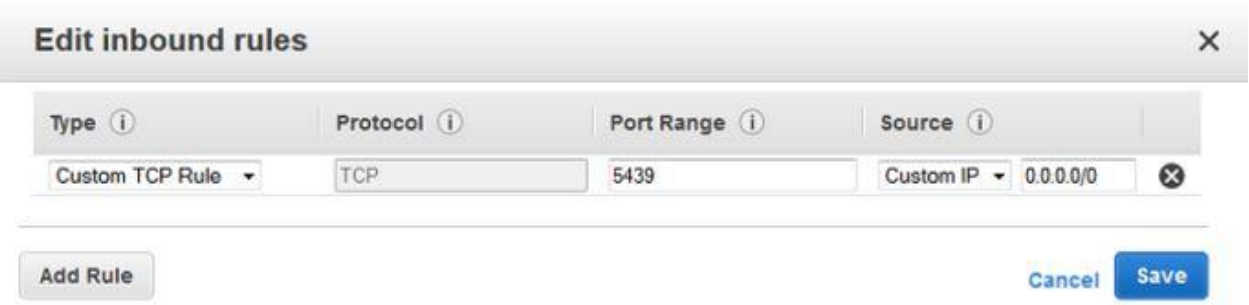
Edit inbound rules

Type	Protocol	Port Range	Source	
Custom TCP Rule	TCP	5439	Custom IP	0.0.0.0/0

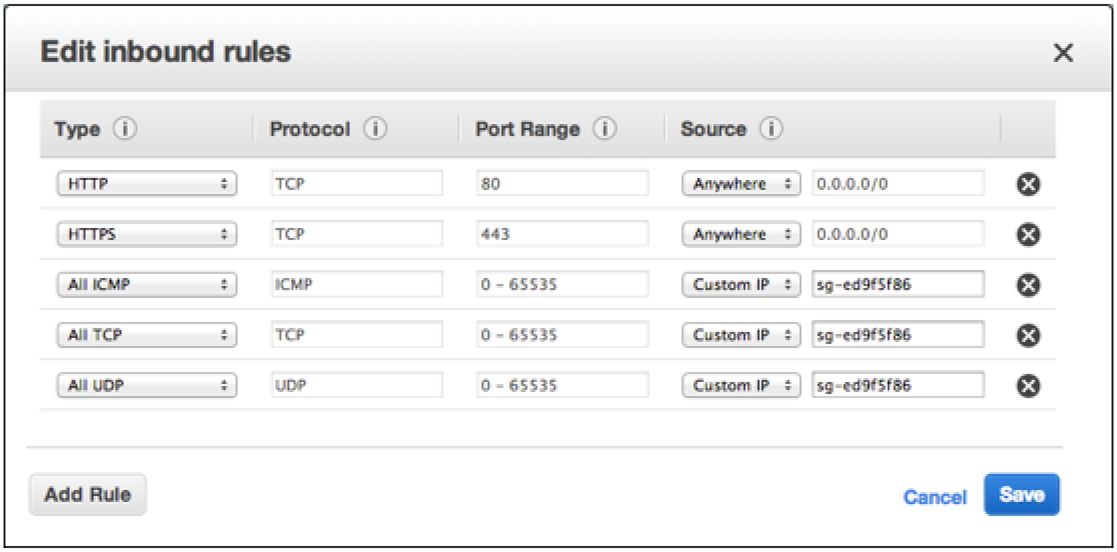

Edit inbound rules

Type	Protocol	Port Range	Source	
HTTP	TCP	80	Anywhere	0.0.0.0/0
HTTPS	TCP	443	Anywhere	0.0.0.0/0
All ICMP	ICMP	0 - 65535	Custom IP	sg-ed9f5f86
All TCP	TCP	0 - 65535	Custom IP	sg-ed9f5f86
All UDP	UDP	0 - 65535	Custom IP	sg-ed9f5f86

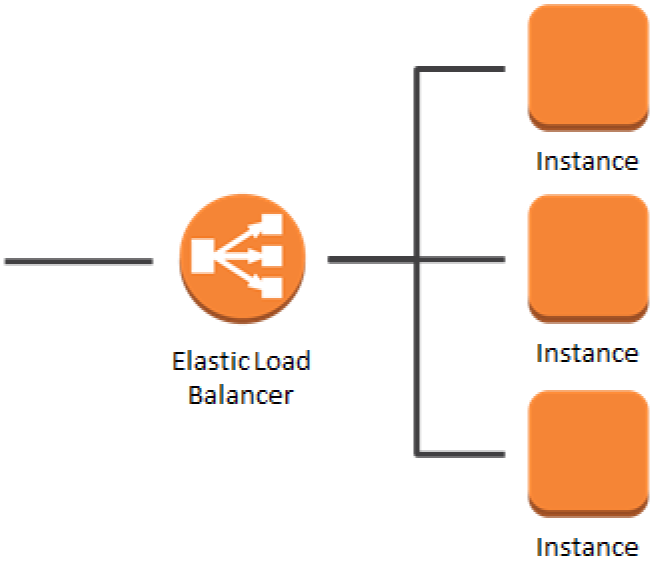

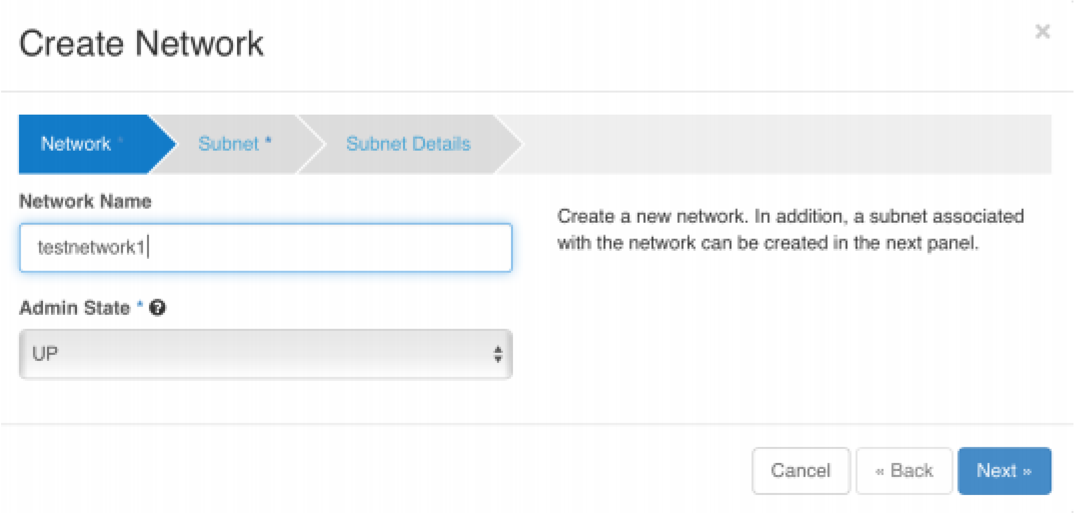

Create Network

Network Name: testnetwork1

Admin State *: UP

Create a new network. In addition, a subnet associated with the network can be created in the next panel.

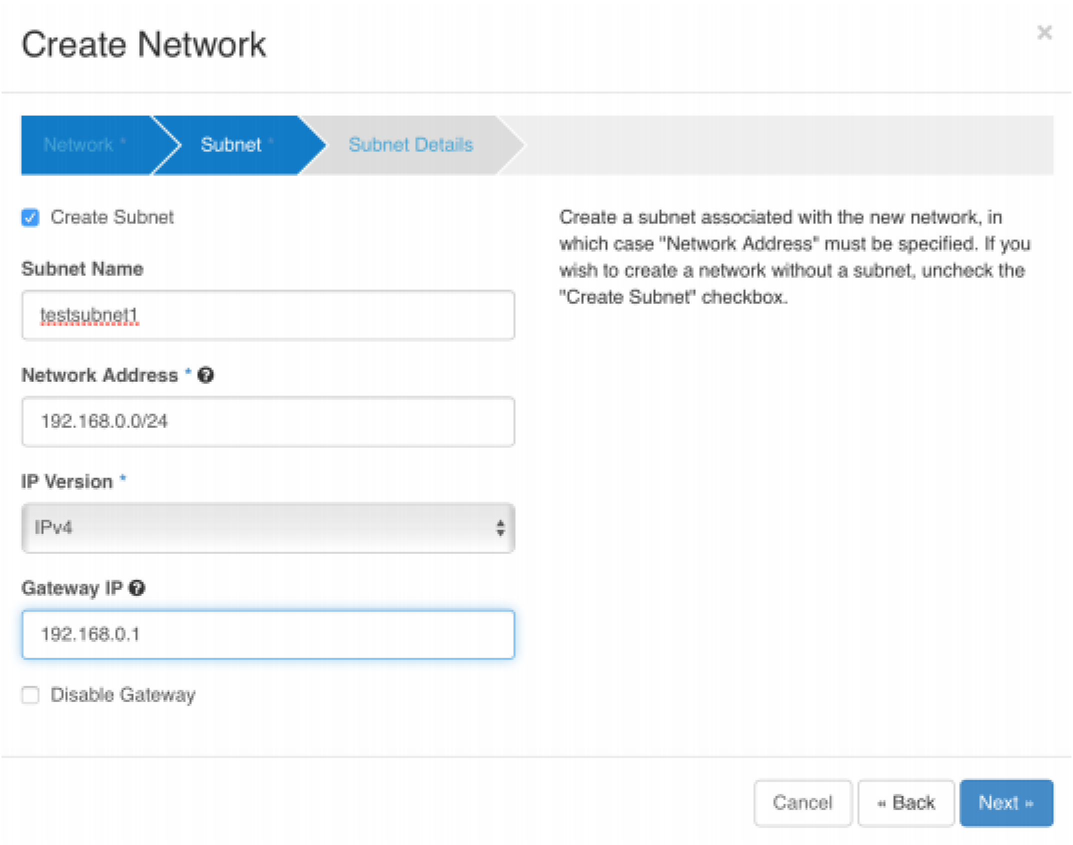

Create Network

☑ Create Subnet

Subnet Name: testsubnet1

Network Address *: 192.168.0.0/24

IP Version *: IPv4

Gateway IP: 192.168.0.1

□ Disable Gateway

Create a subnet associated with the new network, in which case "Network Address" must be specified. If you wish to create a network without a subnet, uncheck the "Create Subnet" checkbox.

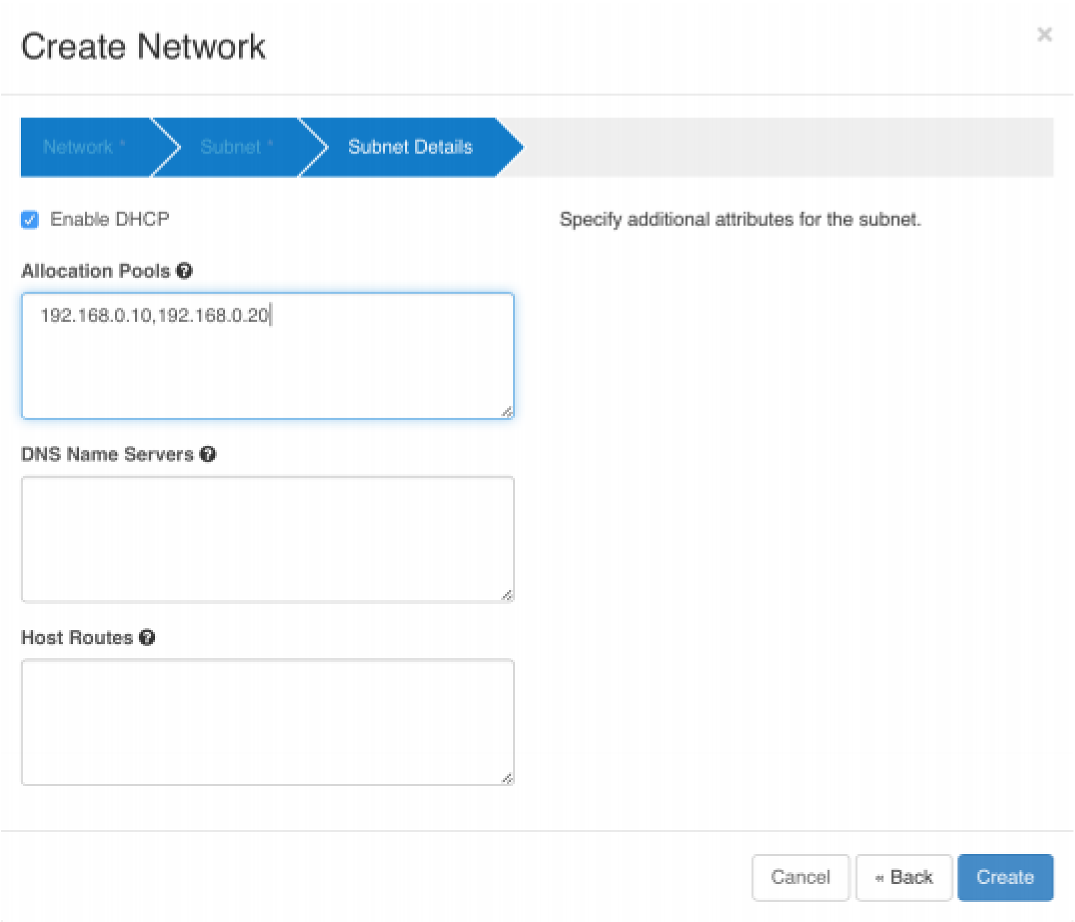

Create Network

☑ Enable DHCP

Specify additional attributes for the subnet.

Allocation Pools

192.168.0.10,192.168.0.20

DNS Name Servers

Host Routes

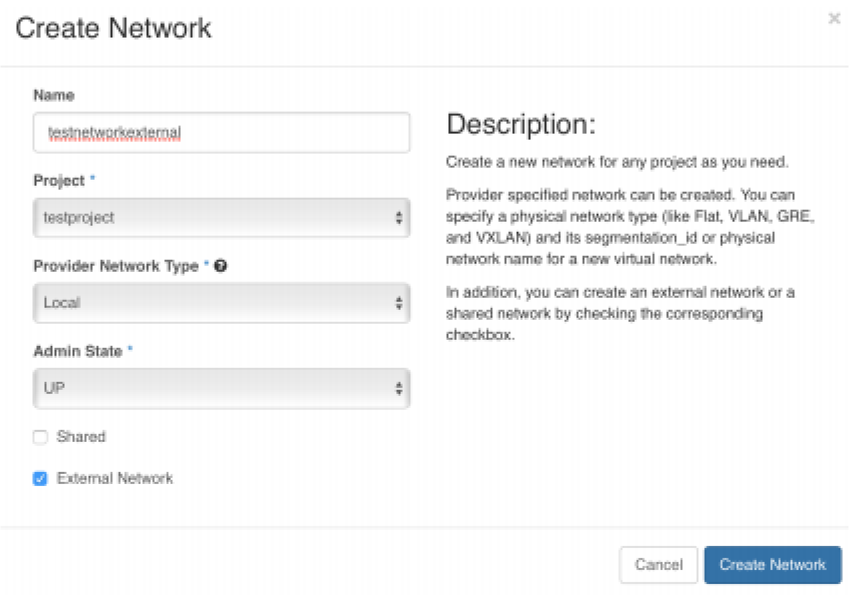

Create Network

Name

testnetworkexternal

Project *

testproject

Provider Network Type *

Local

Admin State *

UP

□ Shared

☑ External Network

Description:

Create a new network for any project as you need.

Provider specified network can be created. You can specify a physical network type (like Flat, VLAN, GRE, and VXLAN) and its segmentation_id or physical network name for a new virtual network.

In addition, you can create an external network or a shared network by checking the corresponding checkbox.

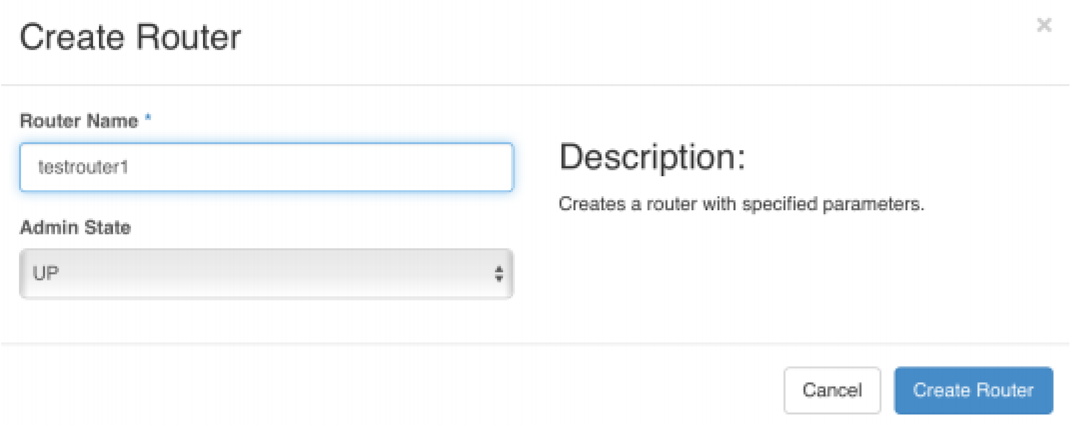

Create Router

Router Name *

testrouter1

Admin State

UP

Description:

Creates a router with specified parameters.

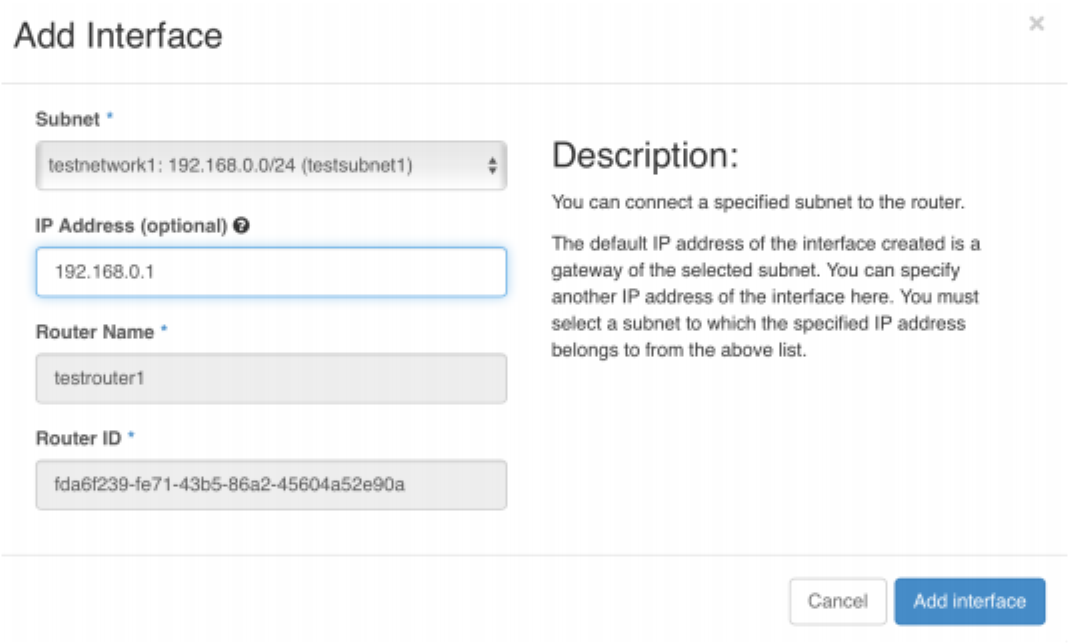

Add Interface

Subnet *

testnetwork1: 192.168.0.0/24 (testsubnet1)

IP Address (optional)

192.168.0.1

Router Name *

testrouter1

Router ID *

fda6f239-fe71-43b5-86a2-45604a52e90a

Description:

You can connect a specified subnet to the router.

The default IP address of the interface created is a gateway of the selected subnet. You can specify another IP address of the interface here. You must select a subnet to which the specified IP address belongs to from the above list.

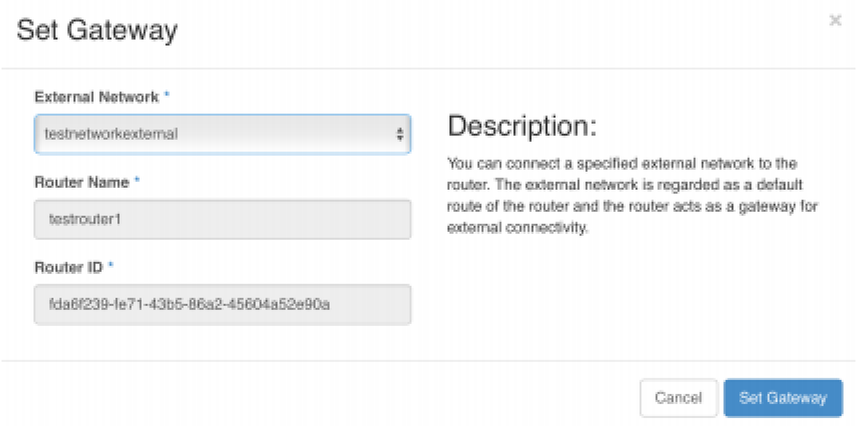

Set Gateway

External Network *

testnetworkexternal

Router Name *

testrouter1

Router ID *

fda6f239-fe71-43b5-86a2-45604a52e90a

Description:

You can connect a specified external network to the router. The external network is regarded as a default route of the router and the router acts as a gateway for external connectivity.

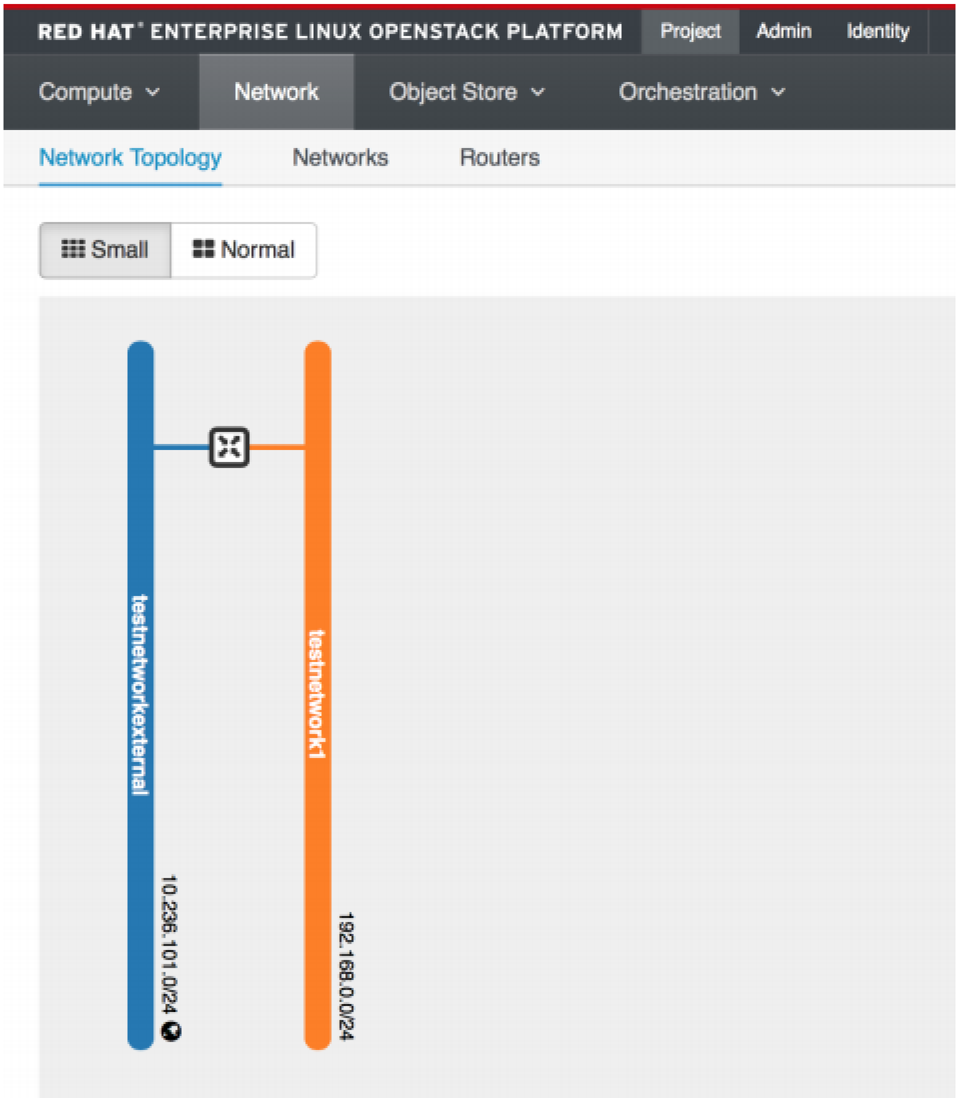

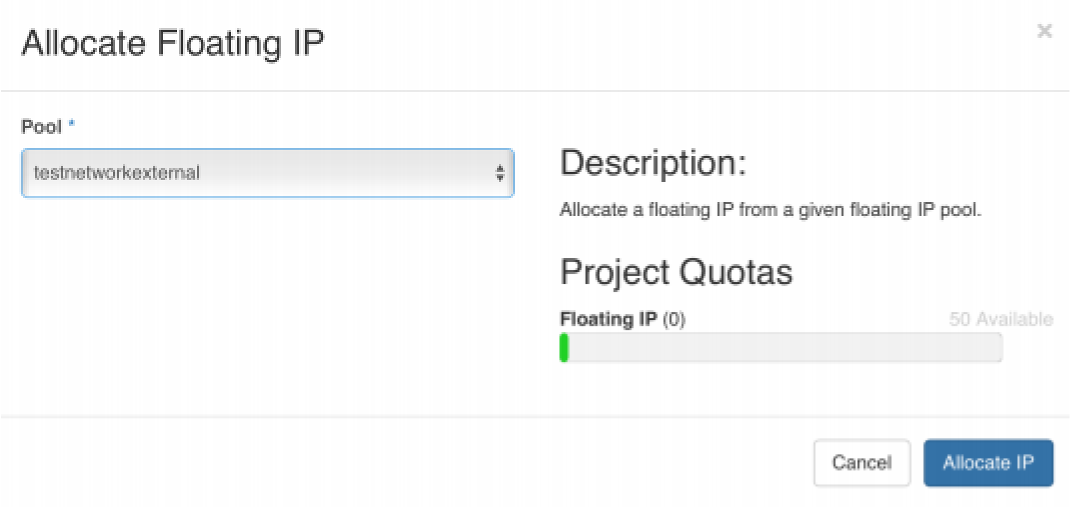
Allocate Floating IP

Pool *

testnetworkexternal

Description:

Allocate a floating IP from a given floating IP pool.

Project Quotas

Floating IP (0) 50 Available

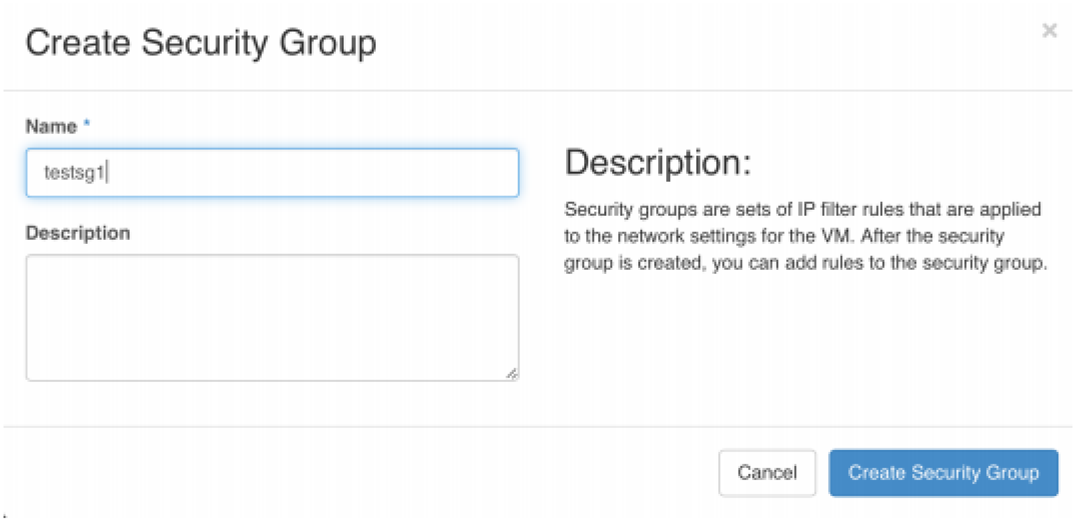

Create Security Group

Name * testsg1

Description

Description:

Security groups are sets of IP filter rules that are applied to the network settings for the VM. After the security group is created, you can add rules to the security group.

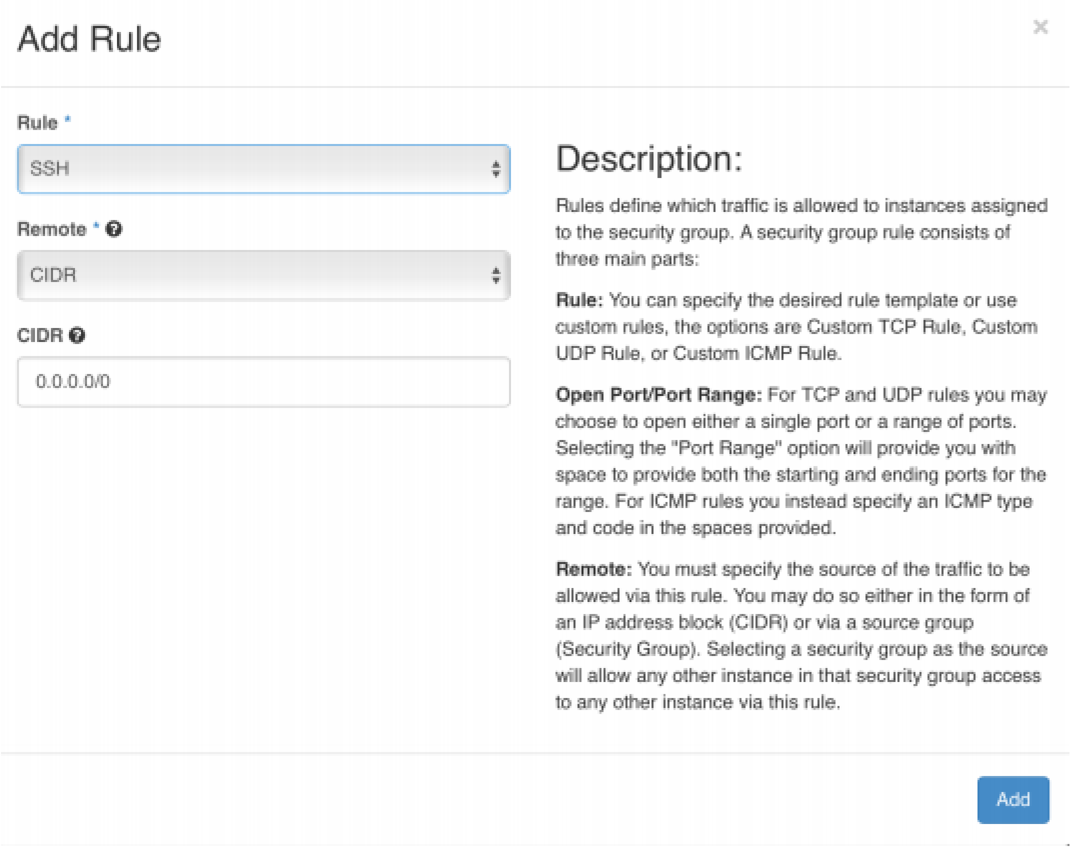

Add Rule

Rule * SSH

Remote * CIDR

CIDR 0.0.0.0/0

Description:

Rules define which traffic is allowed to instances assigned to the security group. A security group rule consists of three main parts:

Rule: You can specify the desired rule template or use custom rules, the options are Custom TCP Rule, Custom UDP Rule, or Custom ICMP Rule.

Open Port/Port Range: For TCP and UDP rules you may choose to open either a single port or a range of ports. Selecting the "Port Range" option will provide you with space to provide both the starting and ending ports for the range. For ICMP rules you instead specify an ICMP type and code in the spaces provided.

Remote: You must specify the source of the traffic to be allowed via this rule. You may do so either in the form of an IP address block (CIDR) or via a source group (Security Group). Selecting a security group as the source will allow any other instance in that security group access to any other instance via this rule.

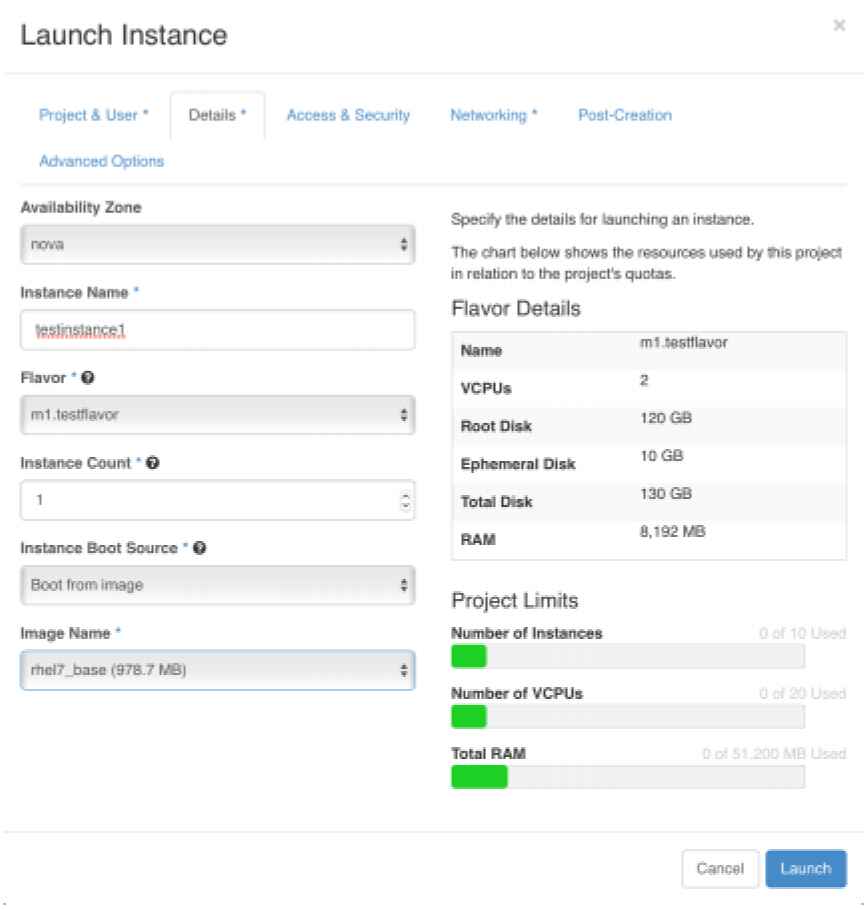

Availability Zone: nova

Instance Name *: testinstance1

Flavor *: m1.testflavor

Instance Count *: 1

Instance Boot Source *: Boot from image

Image Name *: rhel7_base (978.7 MB)

Specify the details for launching an instance.

The chart below shows the resources used by this project in relation to the project's quotas.

Flavor Details

Name	m1.testflavor
VCPUs	2
Root Disk	120 GB
Ephemeral Disk	10 GB
Total Disk	130 GB
RAM	8,192 MB

Project Limits

Number of Instances: 0 of 10 Used

Number of VCPUs: 0 of 20 Used

Total RAM: 0 of 51,200 MB Used

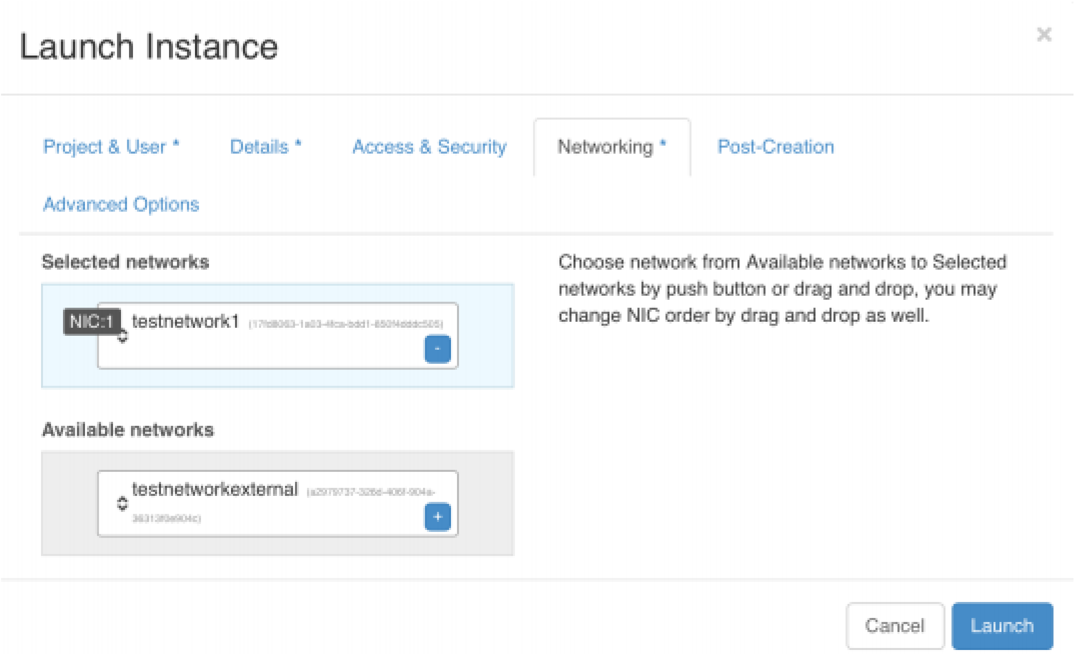

Selected networks

NIC:1 testnetwork1 (17fd8063-1a03-4fca-bdd1-650f4dddc505) -

Available networks

testnetworkexternal (a2979737-326d-406f-904a-363138a904c) +

Choose network from Available networks to Selected networks by push button or drag and drop, you may change NIC order by drag and drop as well.

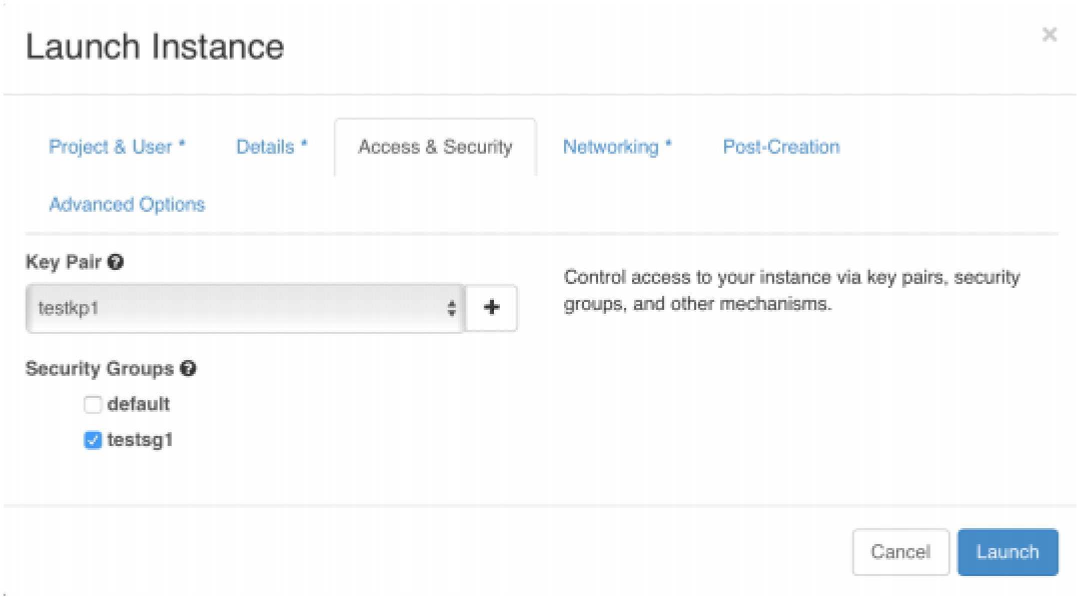

Launch Instance

Key Pair: testkp1

Control access to your instance via key pairs, security groups, and other mechanisms.

Security Groups

□ default

☑ testsg1

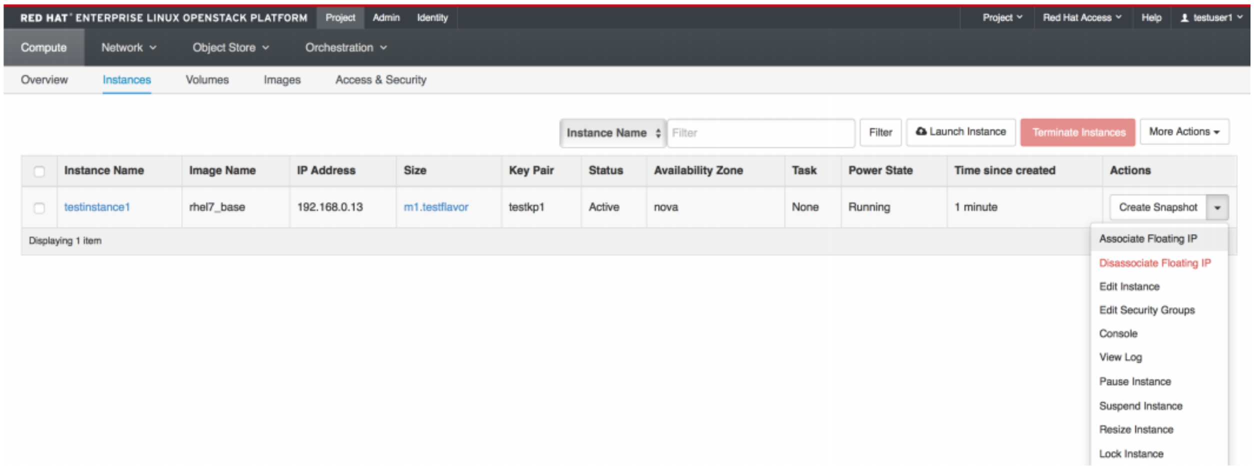

□	Instance Name	Image Name	IP Address	Size	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions
□	testinstance1	rhel7_base	192.168.0.13	m1.testflavor	testkp1	Active	nova	None	Running	1 minute	Create Snapshot

Displaying 1 item

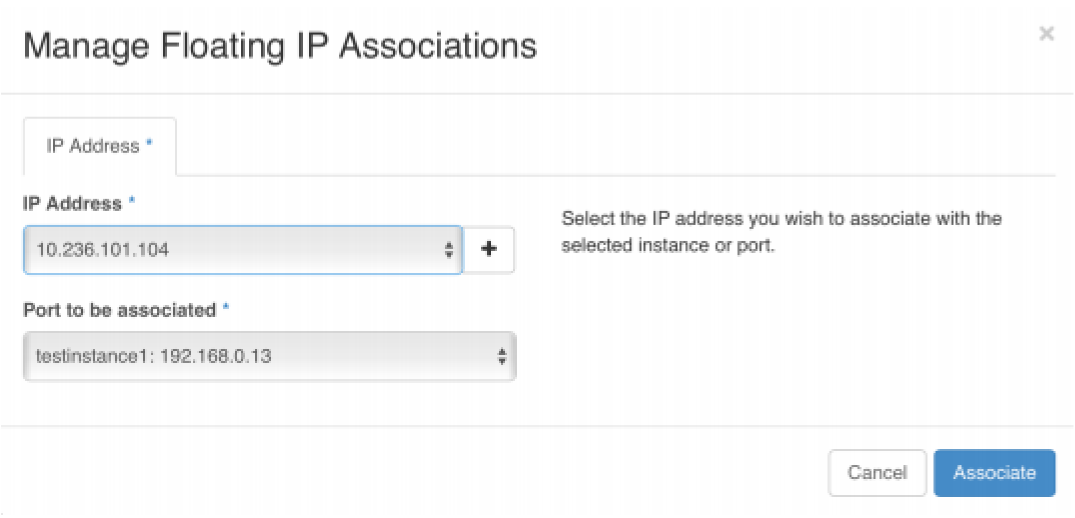

Manage Floating IP Associations

IP Address *: 10.236.101.104

Port to be associated *: testinstance1: 192.168.0.13

Select the IP address you wish to associate with the selected instance or port.

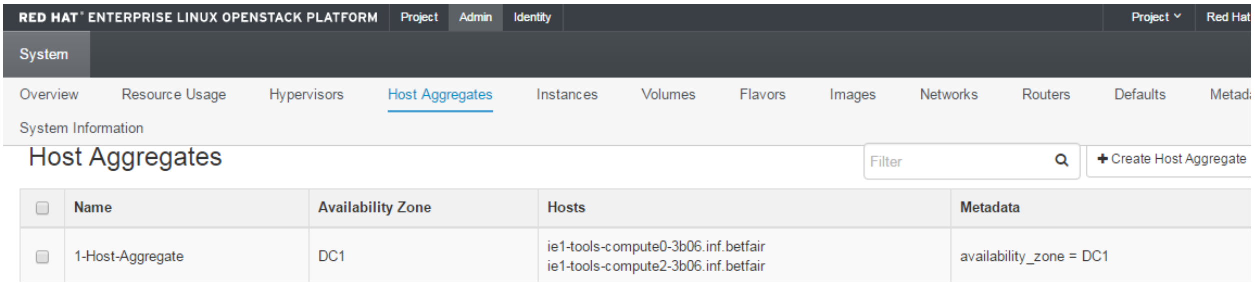

Host Aggregates

□	Name	Availability Zone	Hosts	Metadata
□	1-Host-Aggregate	DC1	ie1-tools-compute0-3b06.inf.betfair ie1-tools-compute2-3b06.inf.betfair	availability_zone = DC1

Chapter 2: The Emergence of Software Defined Networking

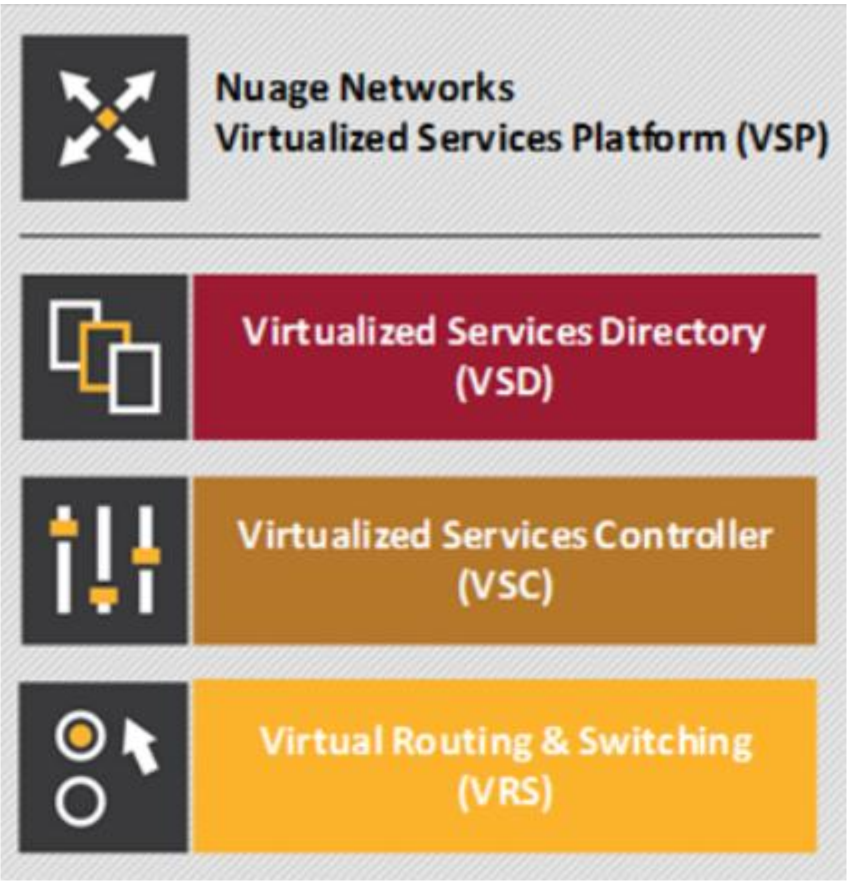

Nuage Networks
Virtualized Services Platform (VSP)

Virtualized Services Directory (VSD)

Virtualized Services Controller (VSC)

Virtual Routing & Switching (VRS)

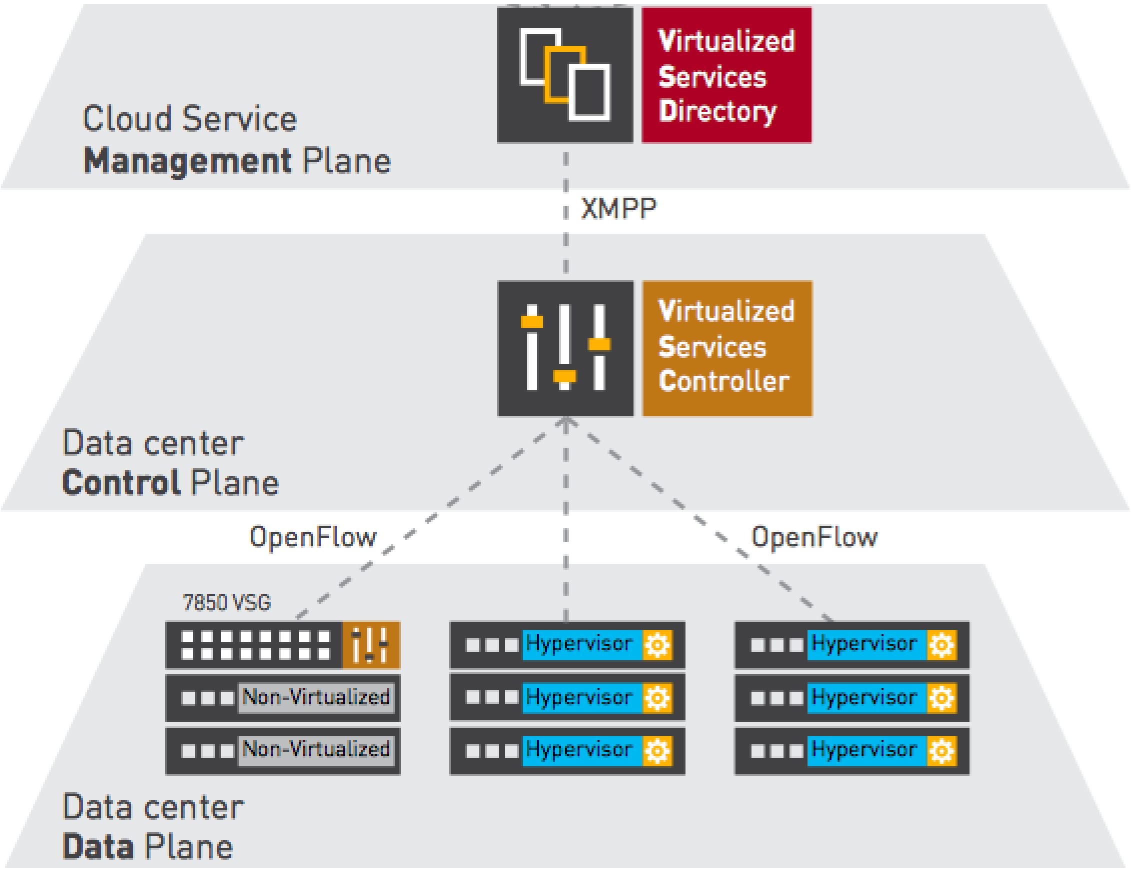
Virtualized Services Directory
Cloud Service Management Plane
XMPP
Virtualized Services Controller
Data center Control Plane
OpenFlow
OpenFlow
7850 VSG
Hypervisor
Hypervisor
Non-Virtualized
Hypervisor
Hypervisor
Non-Virtualized
Hypervisor
Hypervisor
Data center Data Plane

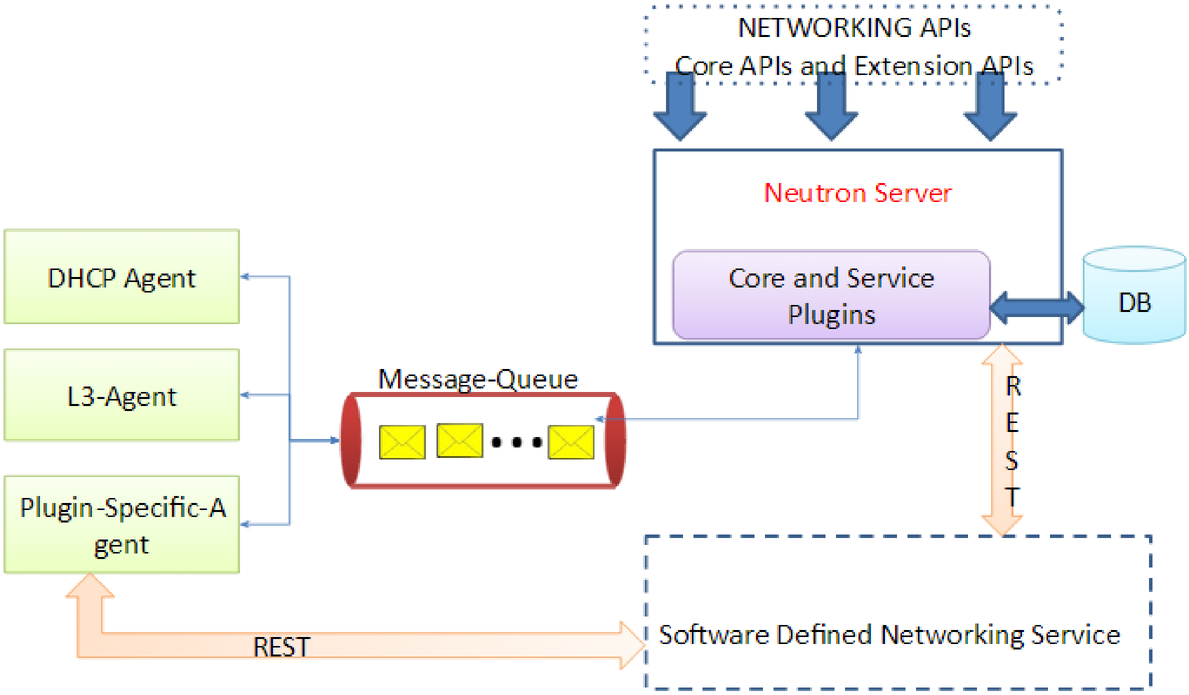
NETWORKING APIs
Core APIs and Extension APIs
Neutron Server
Core and Service Plugins
DB
DHCP Agent
L3-Agent
Plugin-Specific-Agent
Message-Queue
REST
REST
Software Defined Networking Service

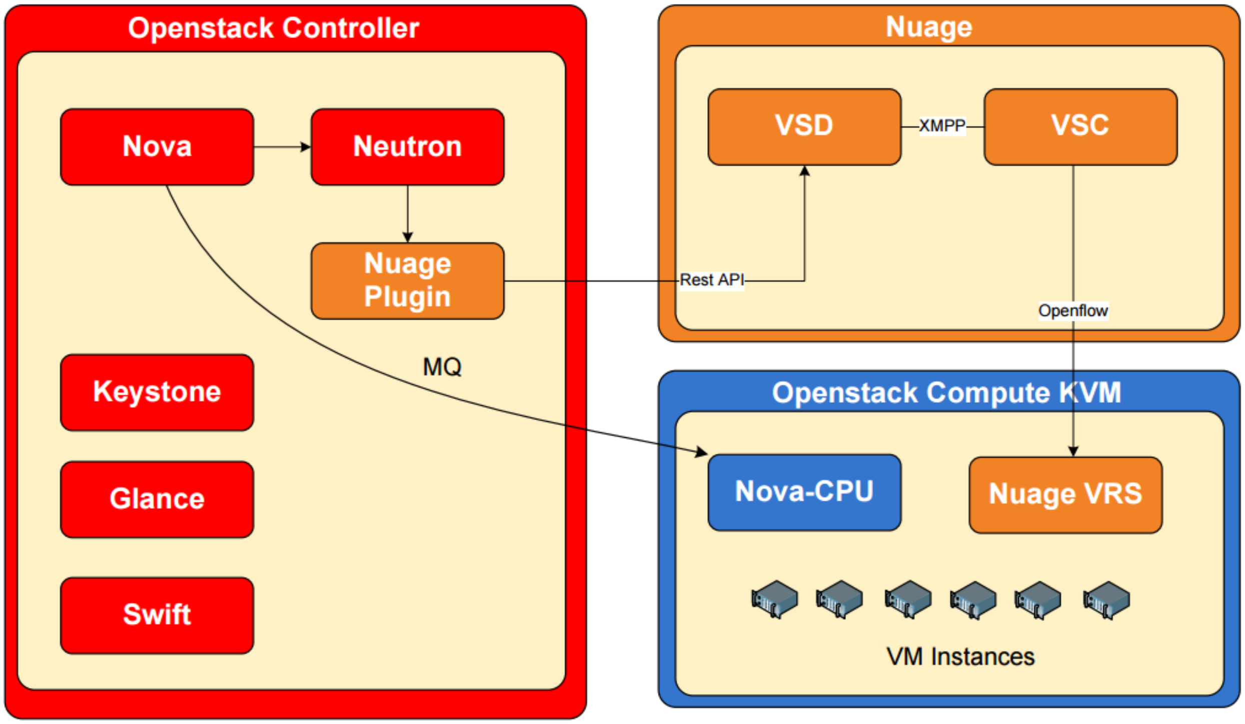

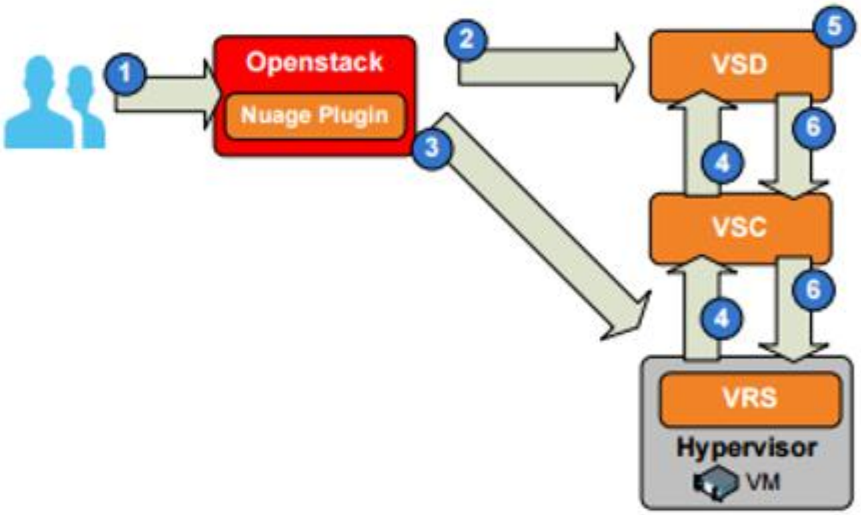

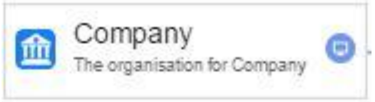
Company
The organisation for Company

Company
The organisation for Company

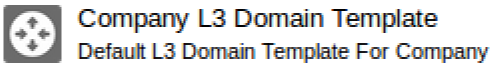
Company L3 Domain Template
Default L3 Domain Template For Company

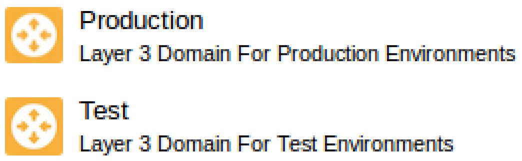

Production
Layer 3 Domain For Production Environments

Test
Layer 3 Domain For Test Environments

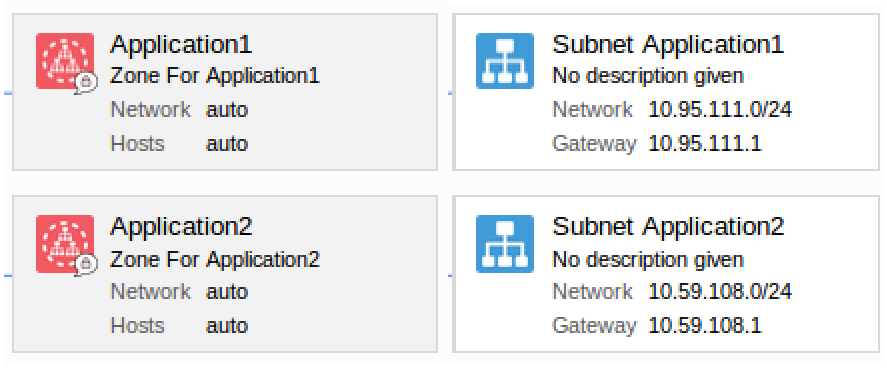

Application1
Zone For Application1

Network	auto
Hosts	auto

Subnet Application1
No description given

Network	10.95.111.0/24
Gateway	10.95.111.1

Application2
Zone For Application2

Network	auto
Hosts	auto

Subnet Application2
No description given

Network	10.59.108.0/24
Gateway	10.59.108.1

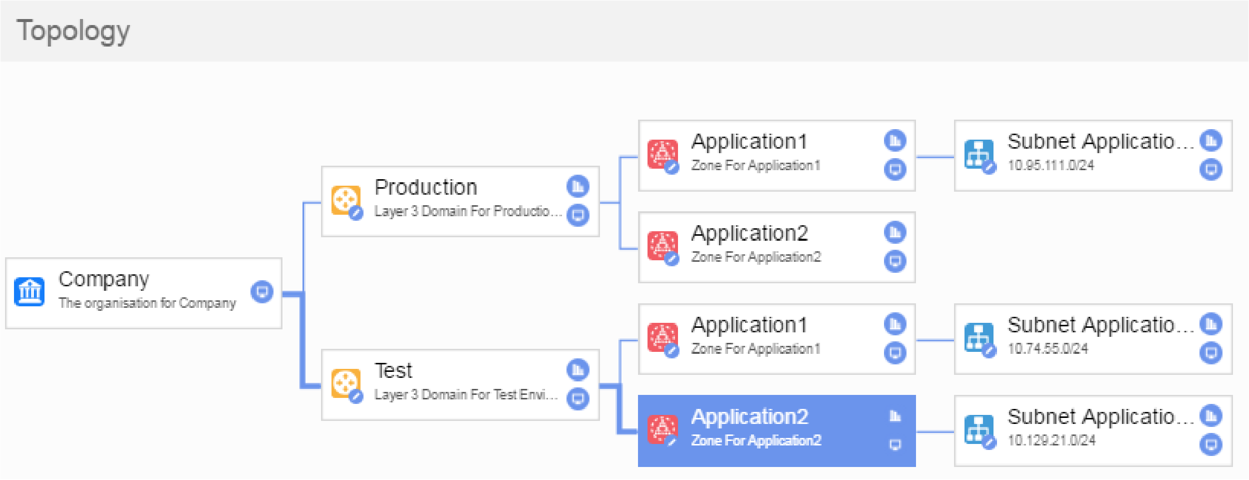

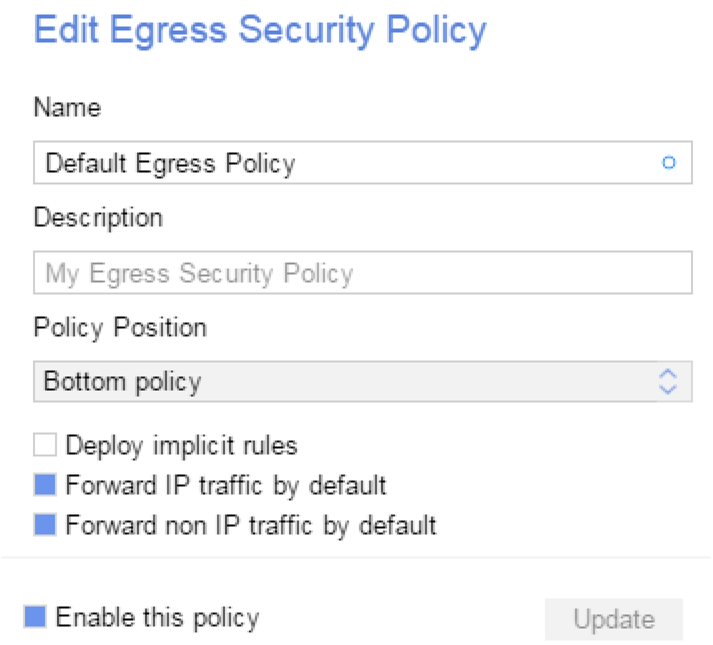

Edit Egress Security Policy

Name

Default Egress Policy

Description

My Egress Security Policy

Policy Position

Bottom policy

☐ Deploy implicit rules

☑ Forward IP traffic by default

☑ Forward non IP traffic by default

☑ Enable this policy

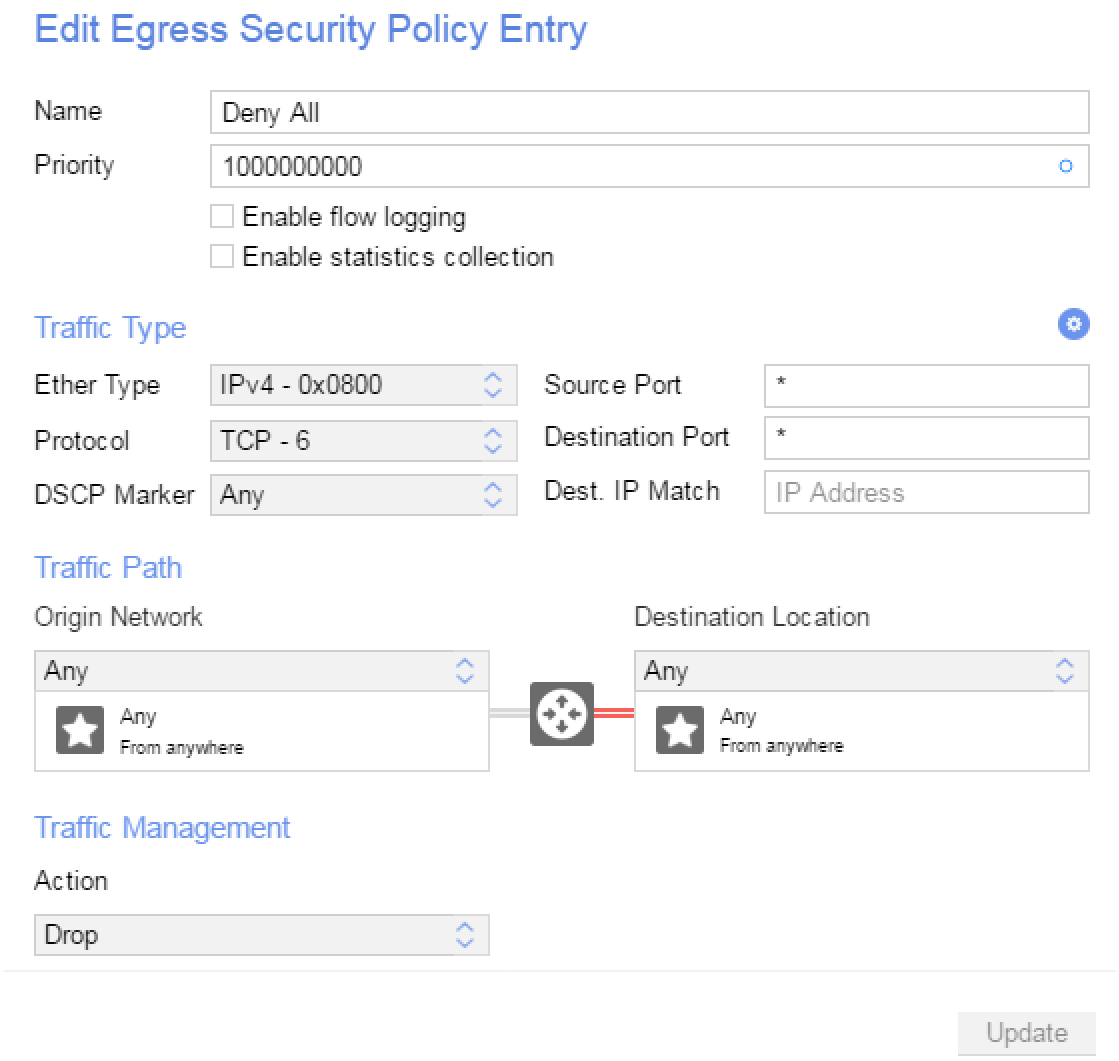

Edit Egress Security Policy Entry

Name: Deny All

Priority: 1000000000

□ Enable flow logging

□ Enable statistics collection

Traffic Type

Ether Type: IPv4 - 0x0800

Protocol: TCP - 6

DSCP Marker: Any

Source Port: *

Destination Port: *

Dest. IP Match: IP Address

Traffic Path

Origin Network: Any

Any
From anywhere

Destination Location: Any

Any
From anywhere

Traffic Management

Action: Drop

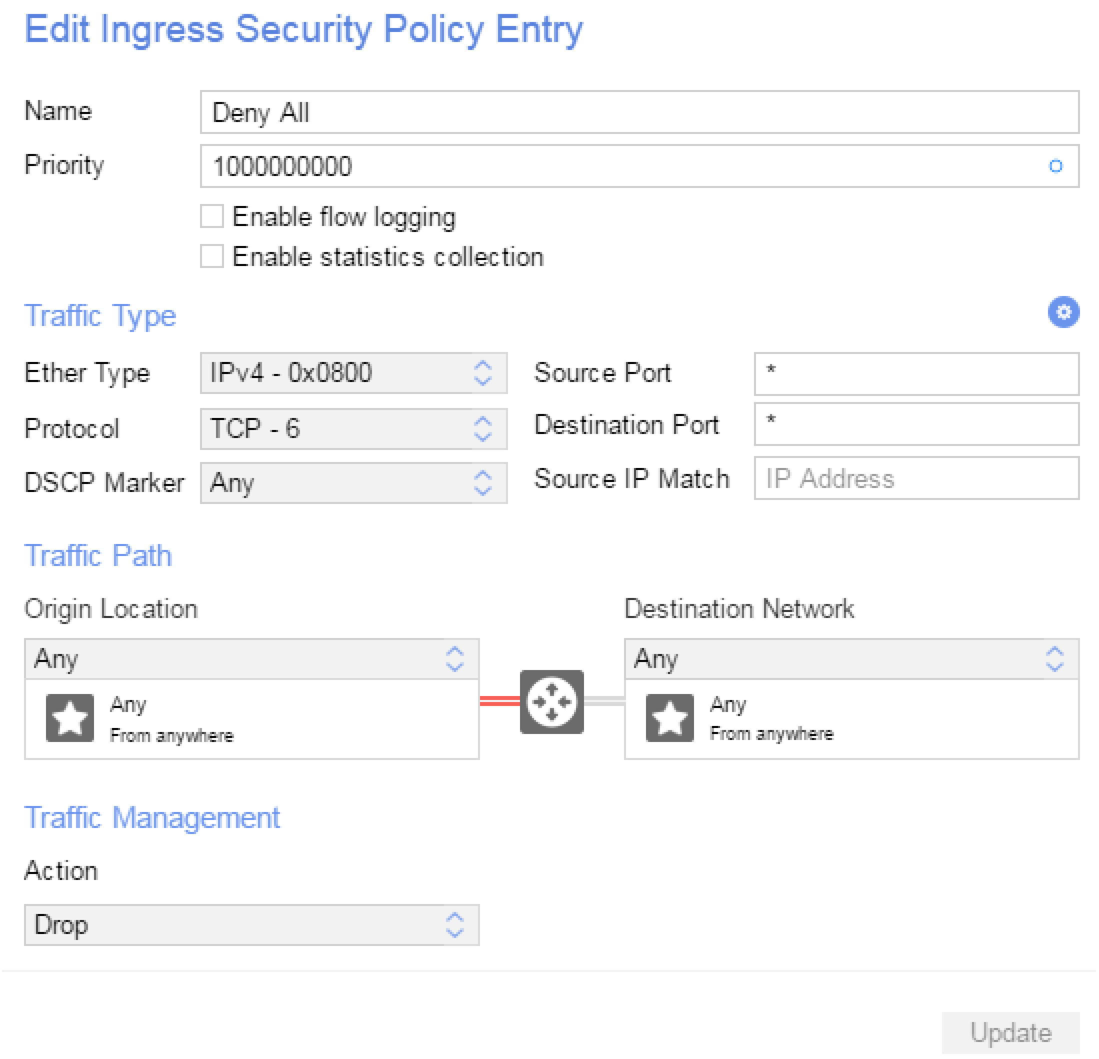

Edit Ingress Security Policy Entry

Name: Deny All

Priority: 1000000000

□ Enable flow logging

□ Enable statistics collection

Traffic Type

Ether Type: IPv4 - 0x0800

Protocol: TCP - 6

DSCP Marker: Any

Source Port: *

Destination Port: *

Source IP Match: IP Address

Traffic Path

Origin Location: Any

Any
From anywhere

Destination Network: Any

Any
From anywhere

Traffic Management

Action: Drop

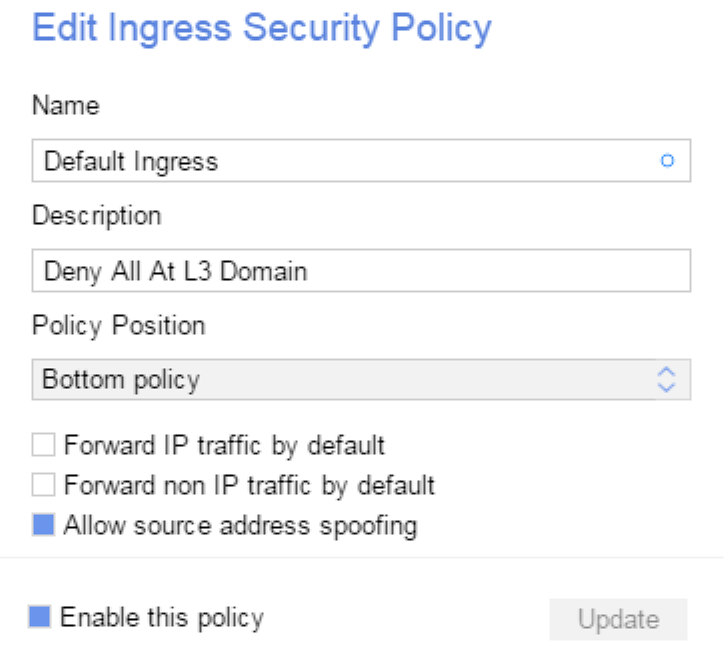

Edit Ingress Security Policy

Name

Default Ingress

Description

Deny All At L3 Domain

Policy Position

Bottom policy

□ Forward IP traffic by default

□ Forward non IP traffic by default

☑ Allow source address spoofing

☑ Enable this policy

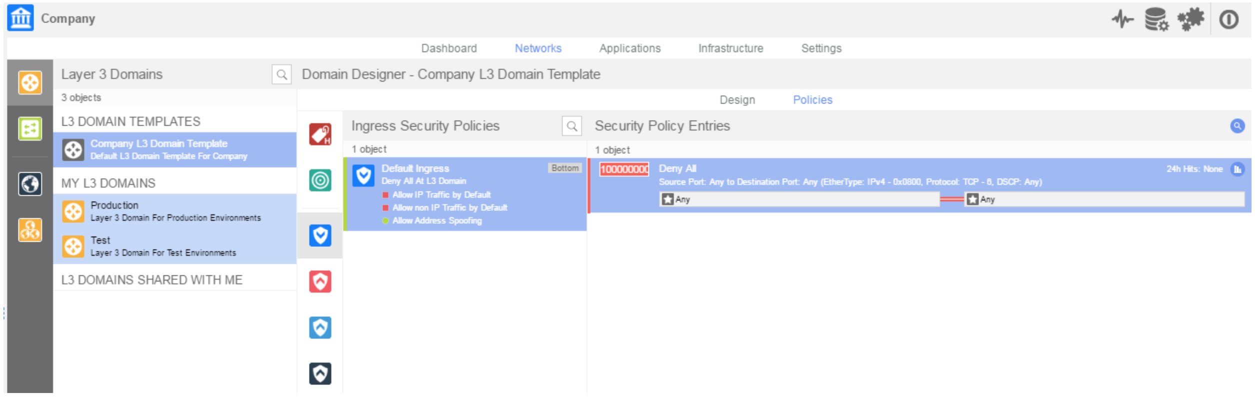

Layer 3 Domains

3 objects

L3 DOMAIN TEMPLATES

Company L3 Domain Template
Default L3 Domain Template For Company

MY L3 DOMAINS

Production
Layer 3 Domain For Production Environments

Test
Layer 3 Domain For Test Environments

L3 DOMAINS SHARED WITH ME

Domain Designer - Company L3 Domain Template

Ingress Security Policies

1 object

Default Ingress — Bottom
Deny All At L3 Domain
- Allow IP Traffic by Default
- Allow non IP Traffic by Default
- Allow Address Spoofing

Security Policy Entries

1 object

100000000 Deny All — 24h Hits: None
Source Port: Any to Destination Port: Any (EtherType: IPv4 - 0x0800, Protocol: TCP - 6, DSCP: Any)
Any — Any

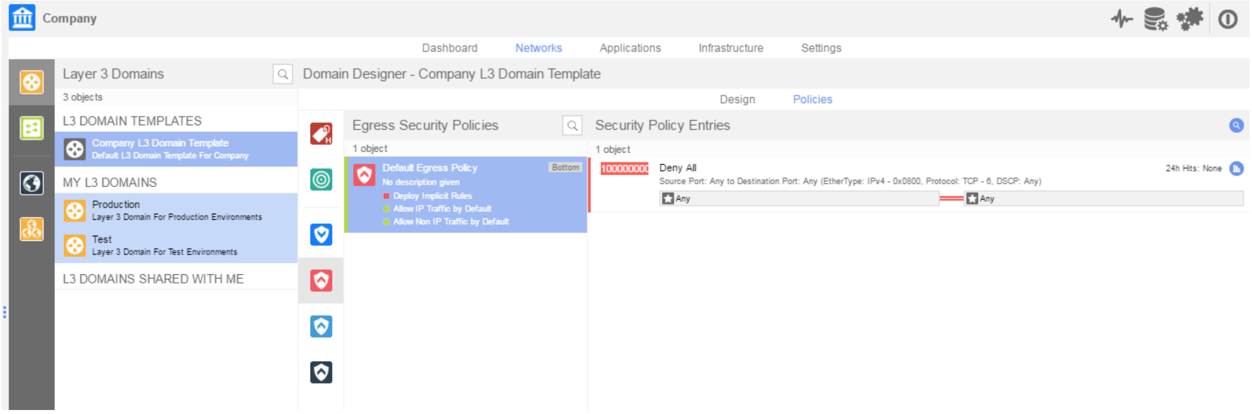

Layer 3 Domains

3 objects

L3 DOMAIN TEMPLATES

Company L3 Domain Template
Default L3 Domain Template For Company

MY L3 DOMAINS

Production
Layer 3 Domain For Production Environments

Test
Layer 3 Domain For Test Environments

L3 DOMAINS SHARED WITH ME

Domain Designer - Company L3 Domain Template

Egress Security Policies

1 object

Default Egress Policy — Bottom
No description given
- Deploy Implicit Rules
- Allow IP Traffic by Default
- Allow Non IP Traffic by Default

Security Policy Entries

1 object

100000000 Deny All — 24h Hits: None
Source Port: Any to Destination Port: Any (EtherType: IPv4 - 0x0800, Protocol: TCP - 6, DSCP: Any)
Any — Any

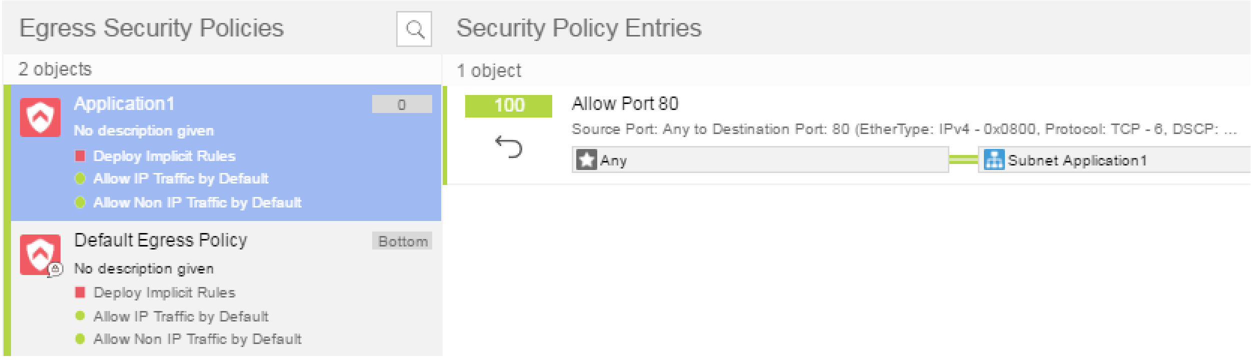

Egress Security Policies

2 objects

Application1 — 0
No description given
- ■ Deploy Implicit Rules
- ● Allow IP Traffic by Default
- ● Allow Non IP Traffic by Default

Default Egress Policy — Bottom
No description given
- ■ Deploy Implicit Rules
- ● Allow IP Traffic by Default
- ● Allow Non IP Traffic by Default

Security Policy Entries

1 object

100 — Allow Port 80
Source Port: Any to Destination Port: 80 (EtherType: IPv4 - 0x0800, Protocol: TCP - 6, DSCP: ...
Any — Subnet Application1

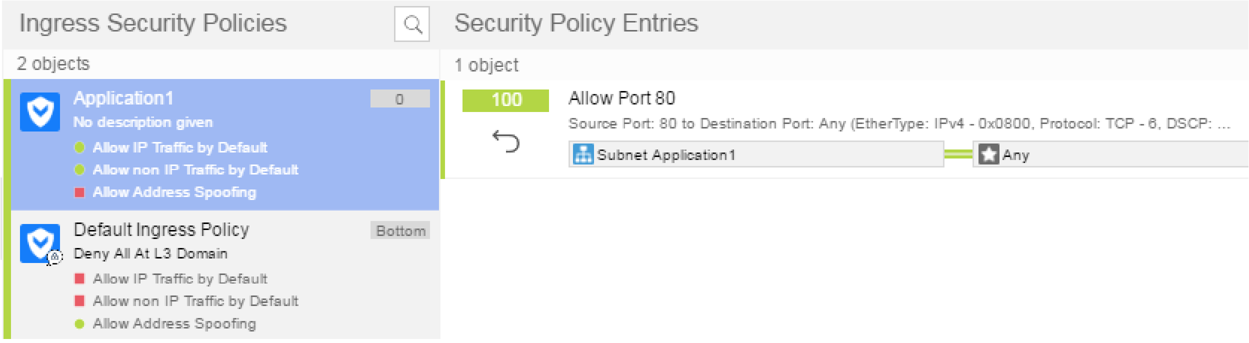

Ingress Security Policies

2 objects

Application1 — 0
No description given
- ● Allow IP Traffic by Default
- ● Allow non IP Traffic by Default
- ■ Allow Address Spoofing

Default Ingress Policy — Bottom
Deny All At L3 Domain
- ■ Allow IP Traffic by Default
- ■ Allow non IP Traffic by Default
- ● Allow Address Spoofing

Security Policy Entries

1 object

100 — Allow Port 80
Source Port: 80 to Destination Port: Any (EtherType: IPv4 - 0x0800, Protocol: TCP - 6, DSCP: ...
Subnet Application1 — Any

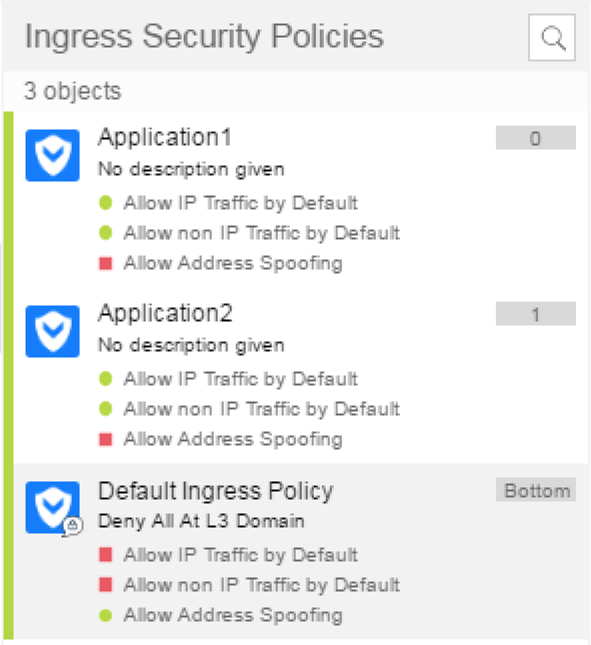

Ingress Security Policies

3 objects

Application1 — 0
No description given
- ● Allow IP Traffic by Default
- ● Allow non IP Traffic by Default
- ■ Allow Address Spoofing

Application2 — 1
No description given
- ● Allow IP Traffic by Default
- ● Allow non IP Traffic by Default
- ■ Allow Address Spoofing

Default Ingress Policy — Bottom
Deny All At L3 Domain
- ■ Allow IP Traffic by Default
- ■ Allow non IP Traffic by Default
- ● Allow Address Spoofing

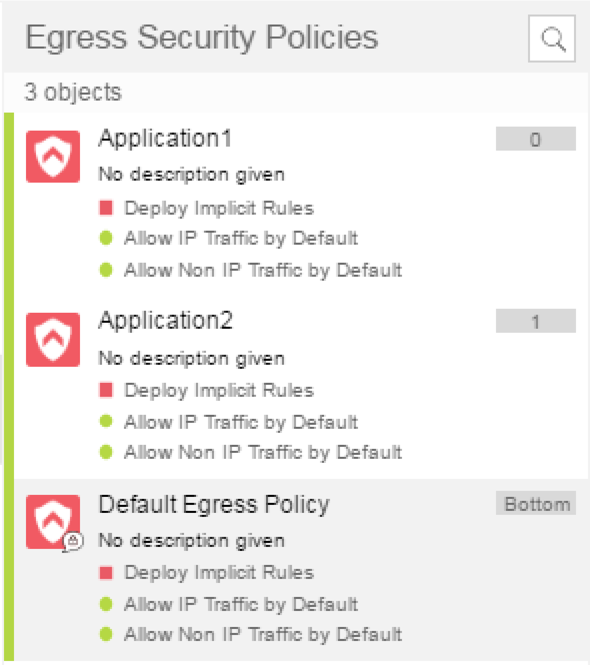

Egress Security Policies

3 objects

Application1 — 0

No description given

- ■ Deploy Implicit Rules
- ● Allow IP Traffic by Default
- ● Allow Non IP Traffic by Default

Application2 — 1

No description given

- ■ Deploy Implicit Rules
- ● Allow IP Traffic by Default
- ● Allow Non IP Traffic by Default

Default Egress Policy — Bottom

No description given

- ■ Deploy Implicit Rules
- ● Allow IP Traffic by Default
- ● Allow Non IP Traffic by Default

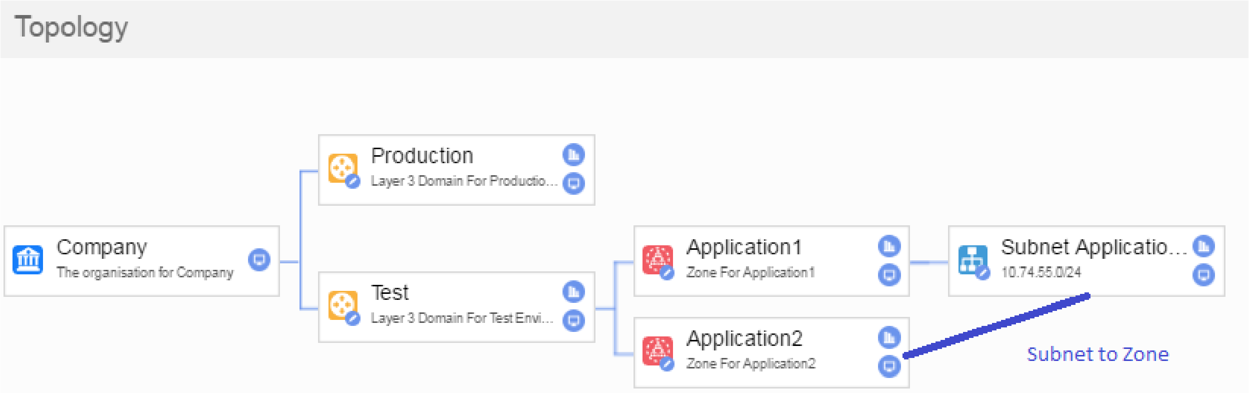

Subnet to Zone

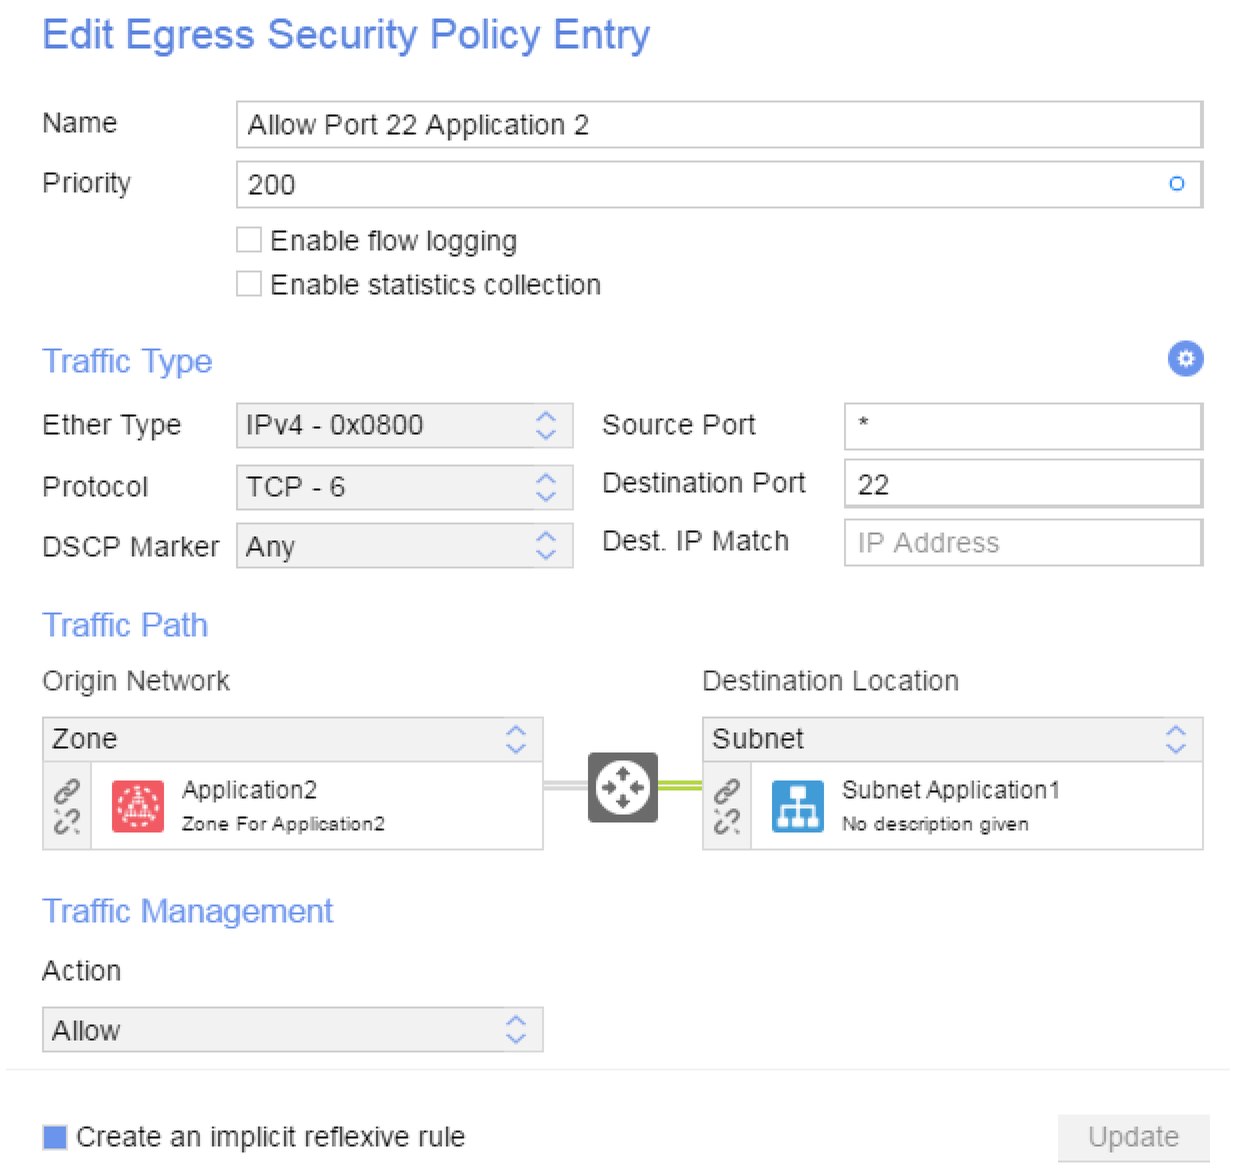

Edit Egress Security Policy Entry

Name: Allow Port 22 Application 2

Priority: 200

☐ Enable flow logging

☐ Enable statistics collection

Traffic Type

Ether Type: IPv4 - 0x0800

Protocol: TCP - 6

DSCP Marker: Any

Source Port: *

Destination Port: 22

Dest. IP Match: IP Address

Traffic Path

Origin Network: Zone

Application2
Zone For Application2

Destination Location: Subnet

Subnet Application1
No description given

Traffic Management

Action: Allow

☑ Create an implicit reflexive rule

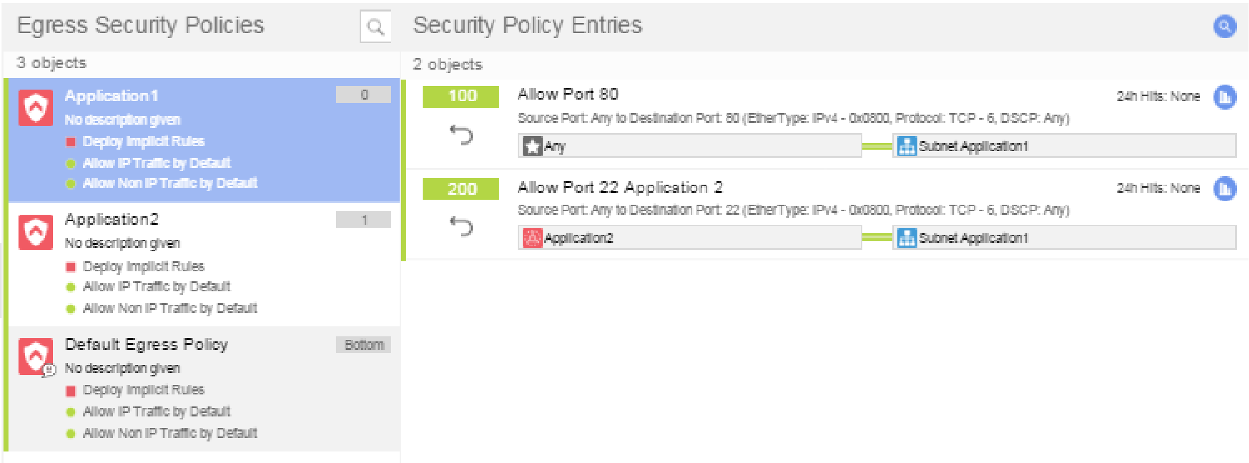

Egress Security Policies

3 objects

Application1 — 0
No description given
- Deploy Implicit Rules
- Allow IP Traffic by Default
- Allow Non IP Traffic by Default

Application2 — 1
No description given
- Deploy Implicit Rules
- Allow IP Traffic by Default
- Allow Non IP Traffic by Default

Default Egress Policy — Bottom
No description given
- Deploy Implicit Rules
- Allow IP Traffic by Default
- Allow Non IP Traffic by Default

Security Policy Entries

2 objects

100 Allow Port 80 — 24h Hits: None
Source Port: Any to Destination Port: 80 (EtherType: IPv4 - 0x0800, Protocol: TCP - 6, DSCP: Any)
Any — Subnet Application1

200 Allow Port 22 Application 2 — 24h Hits: None
Source Port: Any to Destination Port: 22 (EtherType: IPv4 - 0x0800, Protocol: TCP - 6, DSCP: Any)
Application2 — Subnet Application1

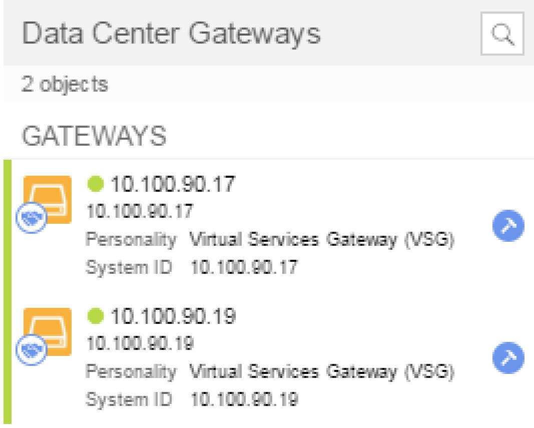

Data Center Gateways

2 objects

GATEWAYS

● 10.100.90.17
10.100.90.17
Personality Virtual Services Gateway (VSG)
System ID 10.100.90.17

● 10.100.90.19
10.100.90.19
Personality Virtual Services Gateway (VSG)
System ID 10.100.90.19

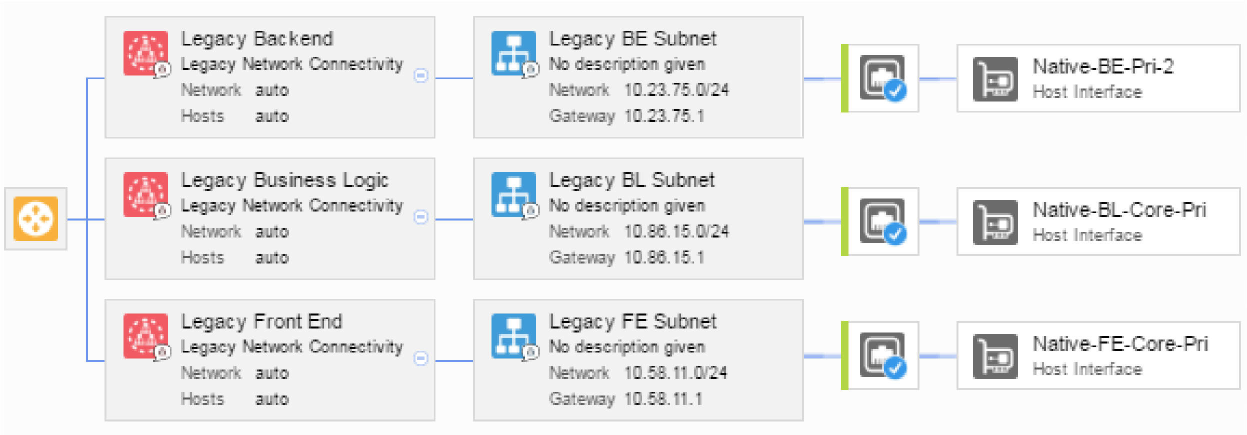

Legacy Backend
Legacy Network Connectivity
Network auto
Hosts auto

Legacy BE Subnet
No description given
Network 10.23.75.0/24
Gateway 10.23.75.1

Native-BE-Pri-2
Host Interface

Legacy Business Logic
Legacy Network Connectivity
Network auto
Hosts auto

Legacy BL Subnet
No description given
Network 10.86.15.0/24
Gateway 10.86.15.1

Native-BL-Core-Pri
Host Interface

Legacy Front End
Legacy Network Connectivity
Network auto
Hosts auto

Legacy FE Subnet
No description given
Network 10.58.11.0/24
Gateway 10.58.11.1

Native-FE-Core-Pri
Host Interface

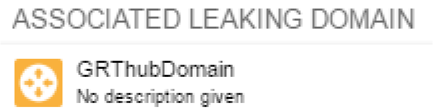

ASSOCIATED LEAKING DOMAIN

GRThubDomain
No description given

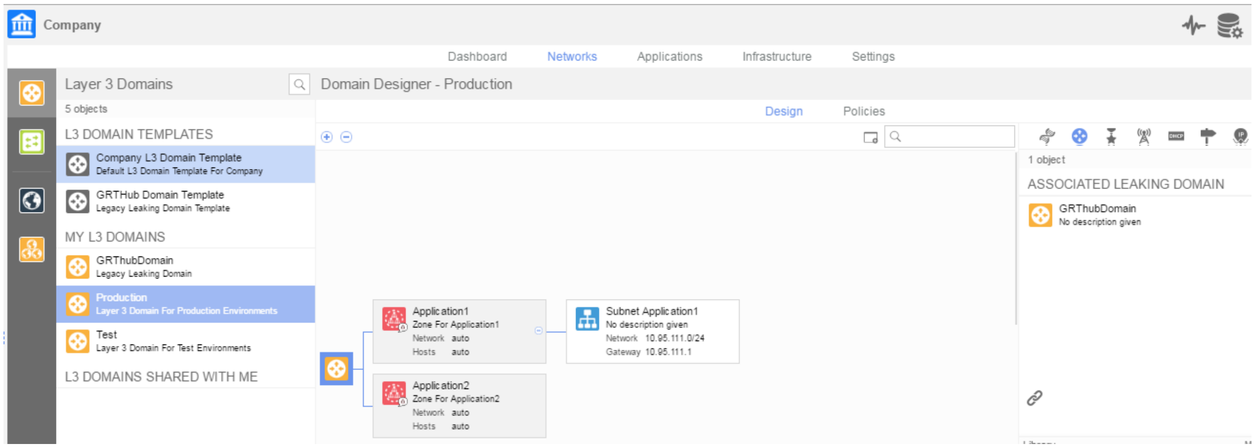

Company

Layer 3 Domains

5 objects

L3 DOMAIN TEMPLATES

Company L3 Domain Template
Default L3 Domain Template For Company

GRTHub Domain Template
Legacy Leaking Domain Template

MY L3 DOMAINS

GRThubDomain
Legacy Leaking Domain

Production
Layer 3 Domain For Production Environments

Test
Layer 3 Domain For Test Environments

L3 DOMAINS SHARED WITH ME

Domain Designer - Production

Application1
Zone For Application1
Network auto
Hosts auto

Subnet Application1
No description given
Network 10.95.111.0/24
Gateway 10.95.111.1

Application2
Zone For Application2
Network auto
Hosts auto

1 object

ASSOCIATED LEAKING DOMAIN

GRThubDomain
No description given

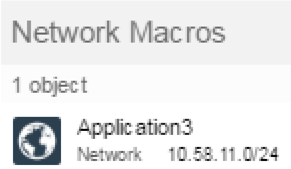

Network Macros

1 object

Application3
Network 10.58.11.0/24

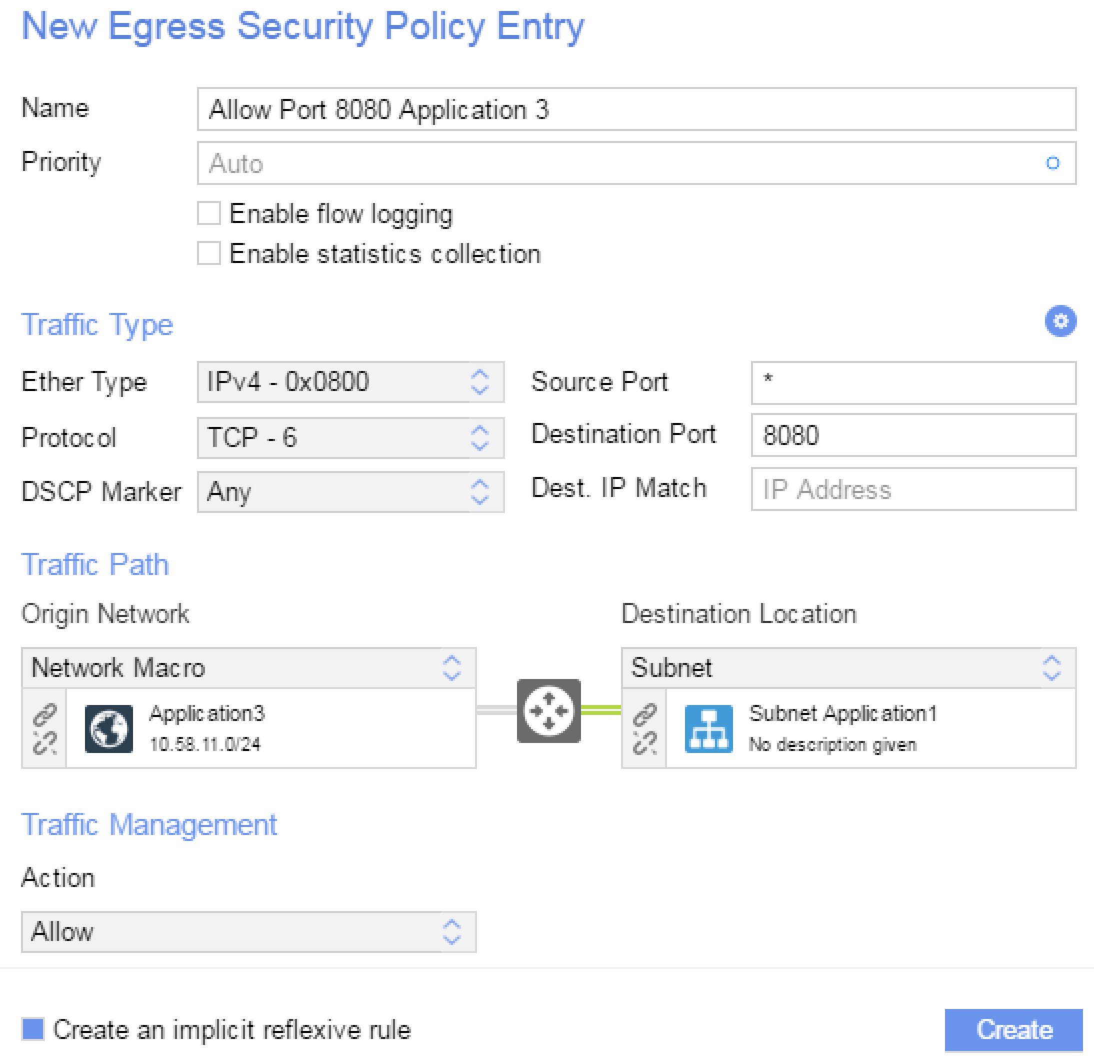

New Egress Security Policy Entry

Name: Allow Port 8080 Application 3

Priority: Auto

☐ Enable flow logging

☐ Enable statistics collection

Traffic Type

Ether Type	IPv4 - 0x0800	Source Port	*
Protocol	TCP - 6	Destination Port	8080
DSCP Marker	Any	Dest. IP Match	IP Address

Traffic Path

Origin Network: Network Macro
Application3
10.58.11.0/24

Destination Location: Subnet
Subnet Application1
No description given

Traffic Management

Action: Allow

☑ Create an implicit reflexive rule

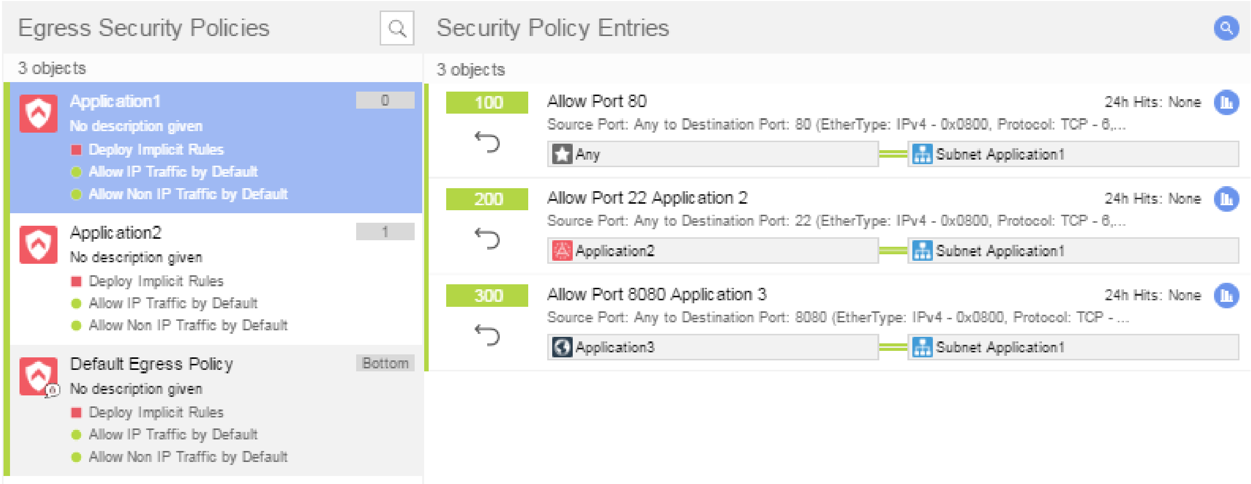

Egress Security Policies

3 objects

Application1 0
No description given
- 🟥 Deploy Implicit Rules
- ● Allow IP Traffic by Default
- ● Allow Non IP Traffic by Default

Application2 1
No description given
- 🟥 Deploy Implicit Rules
- ● Allow IP Traffic by Default
- ● Allow Non IP Traffic by Default

Default Egress Policy Bottom
No description given
- 🟥 Deploy Implicit Rules
- ● Allow IP Traffic by Default
- ● Allow Non IP Traffic by Default

Security Policy Entries

3 objects

100 Allow Port 80 24h Hits: None
Source Port: Any to Destination Port: 80 (EtherType: IPv4 - 0x0800, Protocol: TCP - 6,...
Any = Subnet Application1

200 Allow Port 22 Application 2 24h Hits: None
Source Port: Any to Destination Port: 22 (EtherType: IPv4 - 0x0800, Protocol: TCP - 6,...
Application2 = Subnet Application1

300 Allow Port 8080 Application 3 24h Hits: None
Source Port: Any to Destination Port: 8080 (EtherType: IPv4 - 0x0800, Protocol: TCP - ...
Application3 = Subnet Application1

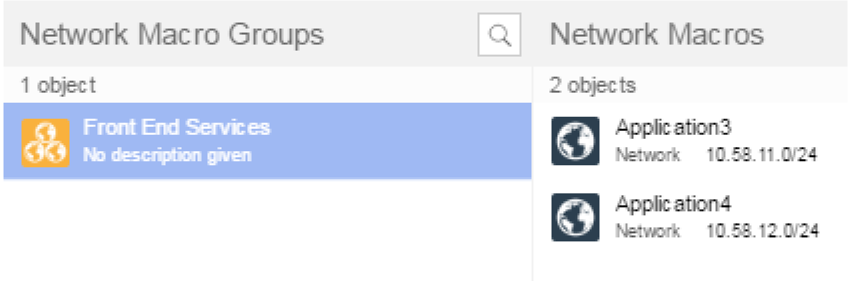

Network Macro Groups

1 object

Front End Services
No description given

Network Macros

2 objects

Application3
Network 10.58.11.0/24

Application4
Network 10.58.12.0/24

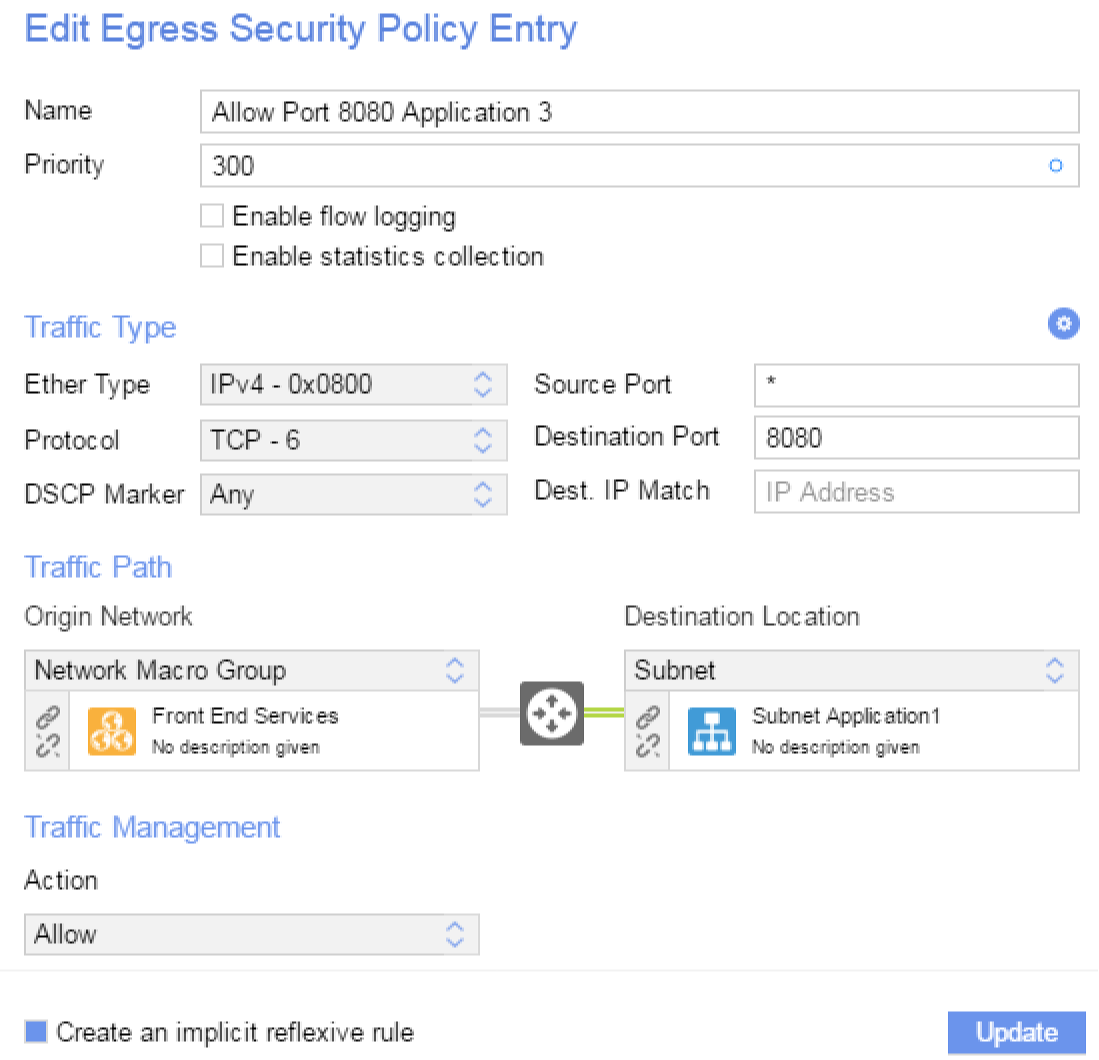

Edit Egress Security Policy Entry

Name	Allow Port 8080 Application 3
Priority	300

□ Enable flow logging

□ Enable statistics collection

Traffic Type

Ether Type	IPv4 - 0x0800	Source Port	*
Protocol	TCP - 6	Destination Port	8080
DSCP Marker	Any	Dest. IP Match	IP Address

Traffic Path

Origin Network	Destination Location
Network Macro Group	Subnet
Front End Services No description given	Subnet Application1 No description given

Traffic Management

Action

Allow

■ Create an implicit reflexive rule

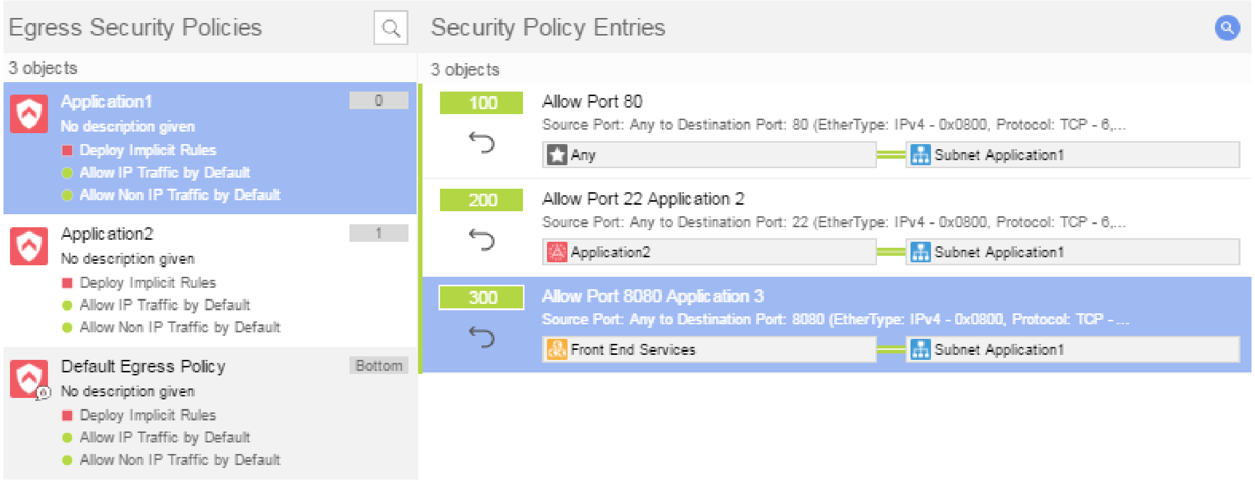

Egress Security Policies

3 objects

- Application1 — 0
  No description given
  - Deploy Implicit Rules
  - Allow IP Traffic by Default
  - Allow Non IP Traffic by Default
- Application2 — 1
  No description given
  - Deploy Implicit Rules
  - Allow IP Traffic by Default
  - Allow Non IP Traffic by Default
- Default Egress Policy — Bottom
  No description given
  - Deploy Implicit Rules
  - Allow IP Traffic by Default
  - Allow Non IP Traffic by Default

Security Policy Entries

3 objects

100	Allow Port 80 Source Port: Any to Destination Port: 80 (EtherType: IPv4 - 0x0800, Protocol: TCP - 6,...	Any	Subnet Application1
200	Allow Port 22 Application 2 Source Port: Any to Destination Port: 22 (EtherType: IPv4 - 0x0800, Protocol: TCP - 6,...	Application2	Subnet Application1
300	Allow Port 8080 Application 3 Source Port: Any to Destination Port: 8080 (EtherType: IPv4 - 0x0800, Protocol: TCP - ...	Front End Services	Subnet Application1

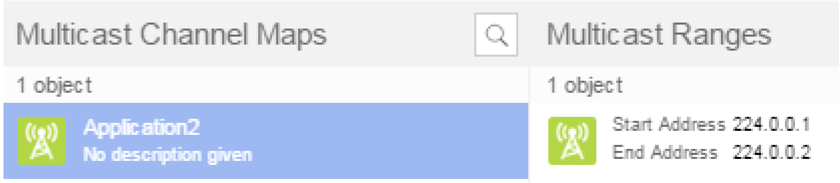

1 object

Application2
No description given

1 object

Start Address 224.0.0.1
End Address 224.0.0.2

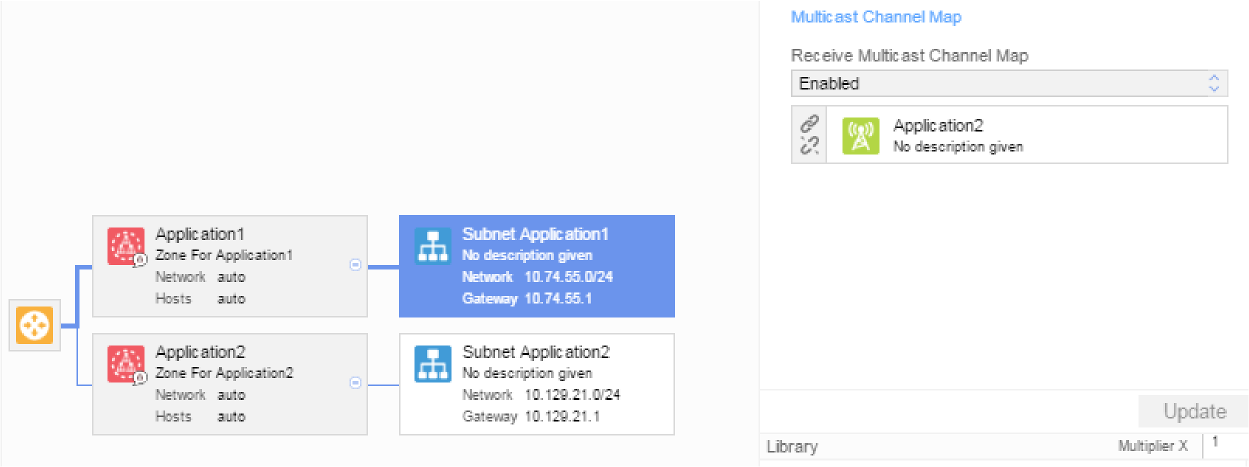

Application1
Zone For Application1
Network auto
Hosts auto

Subnet Application1
No description given
Network 10.74.55.0/24
Gateway 10.74.55.1

Application2
Zone For Application2
Network auto
Hosts auto

Subnet Application2
No description given
Network 10.129.21.0/24
Gateway 10.129.21.1

Receive Multicast Channel Map: Enabled

Application2
No description given

Multiplier X 1

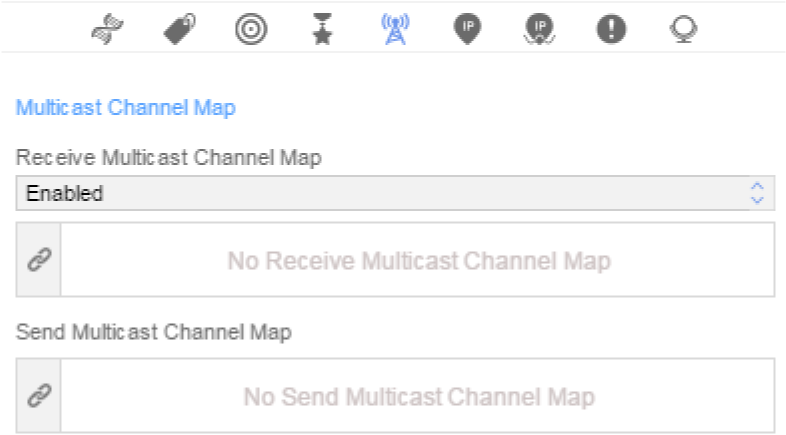

Receive Multicast Channel Map: Enabled

No Receive Multicast Channel Map

Send Multicast Channel Map

No Send Multicast Channel Map

Chapter 3: Bringing DevOps to Network Operations

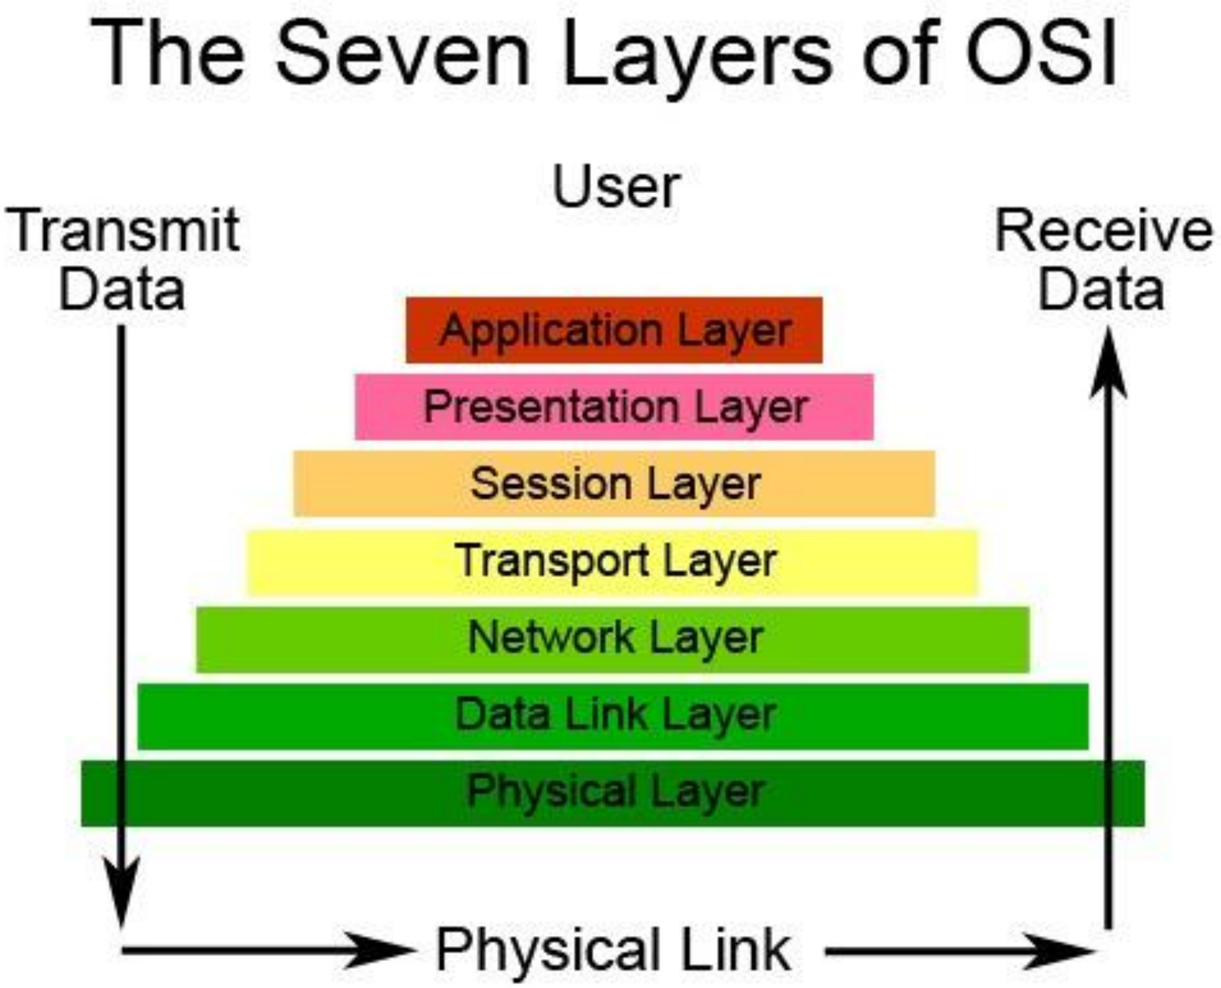

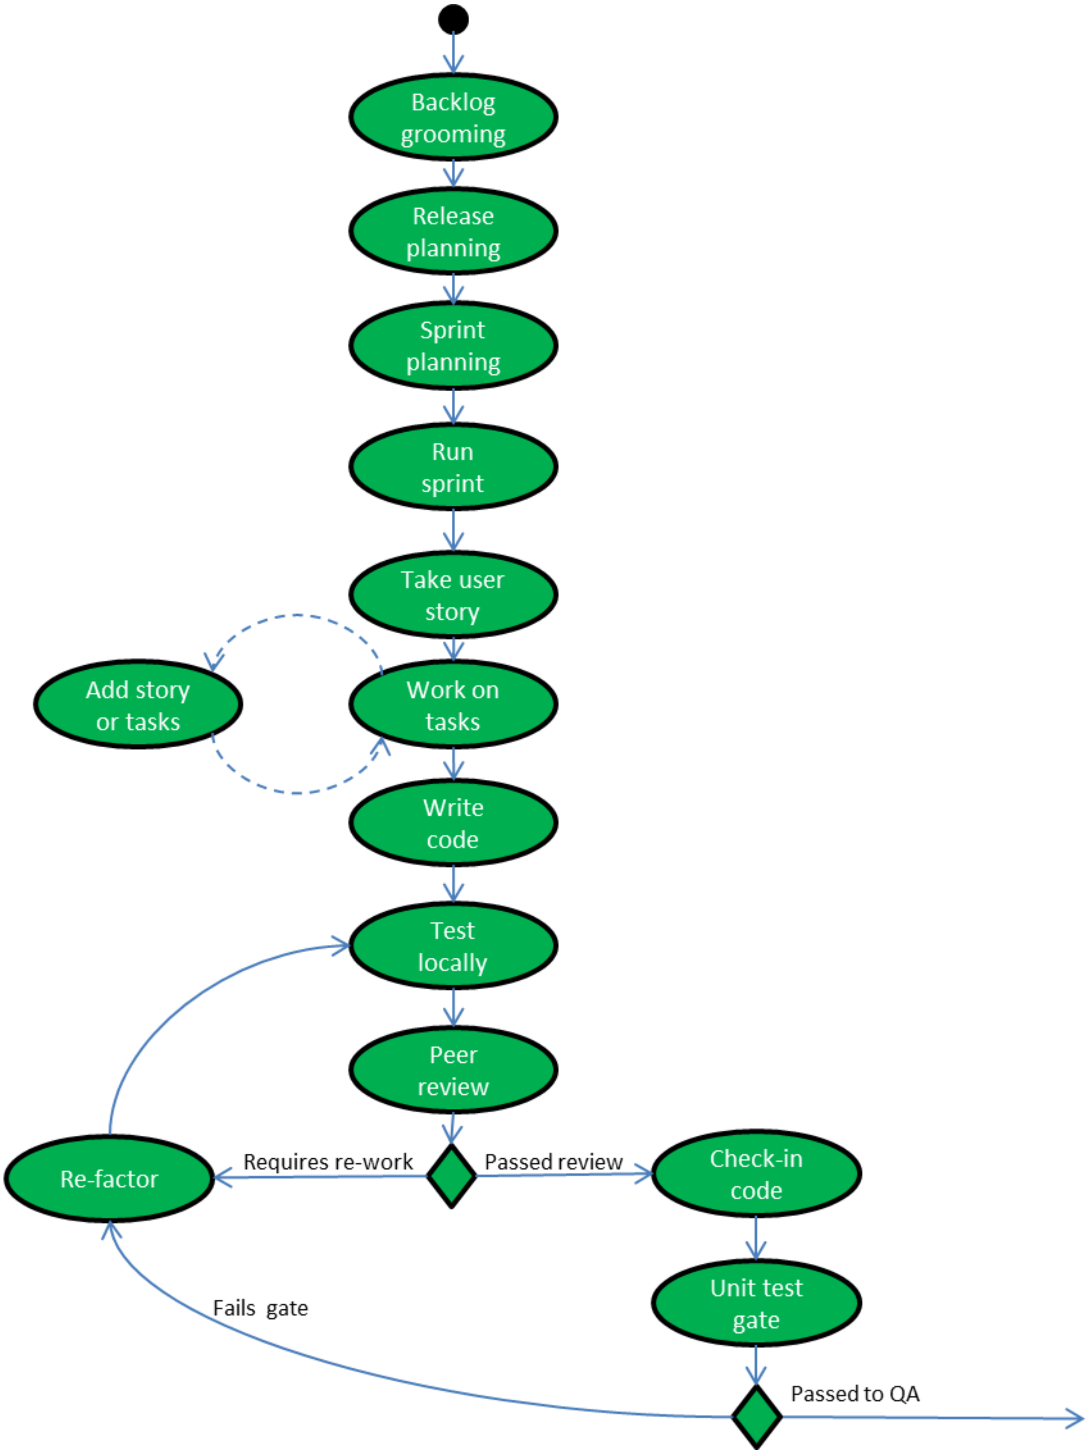
Backlog grooming
Release planning
Sprint planning
Run sprint
Take user story
Add story or tasks
Work on tasks
Write code
Test locally
Peer review
Re-factor
Requires re-work
Passed review
Check-in code
Unit test gate
Fails gate
Passed to QA

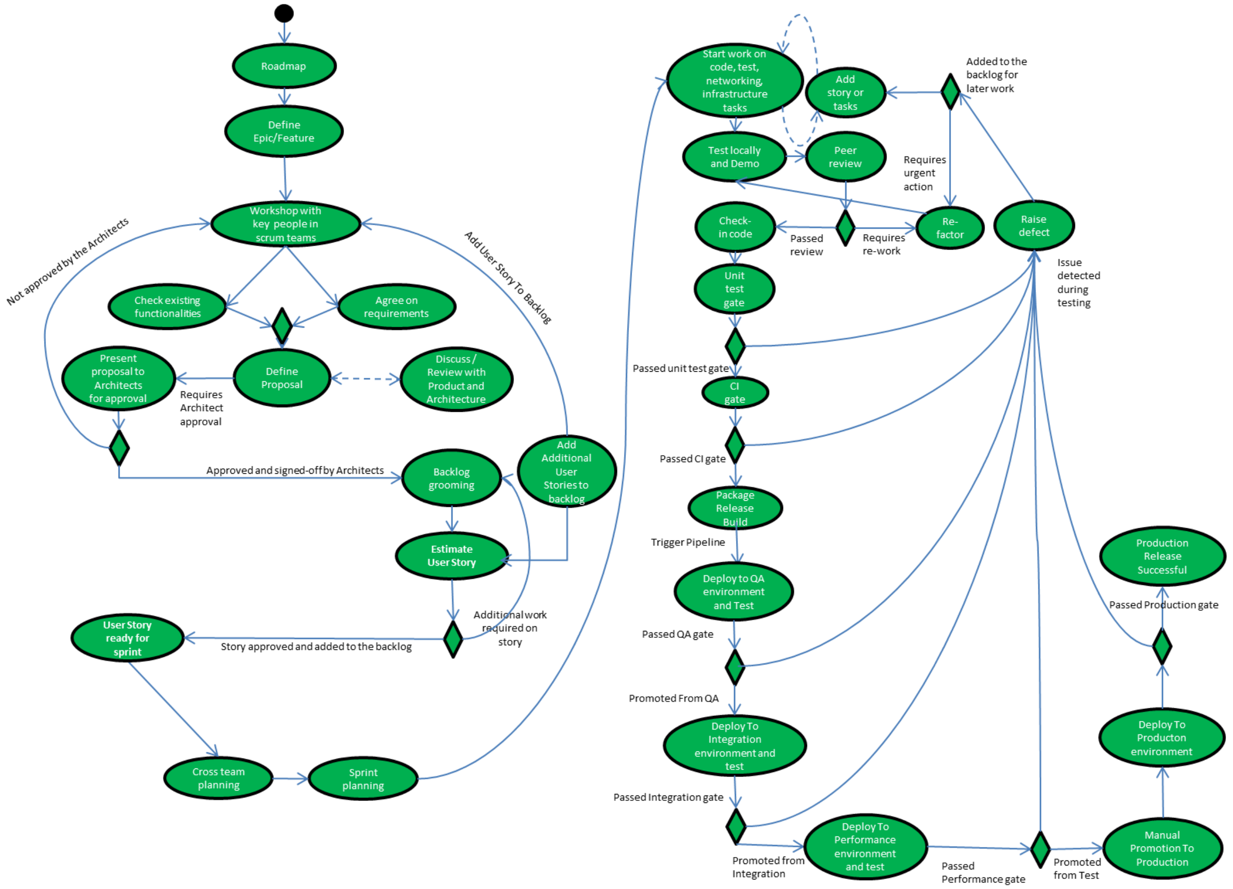
Roadmap
Define Epic/Feature
Workshop with key people in scrum teams
Not approved by the Architects
Add User Story To Backlog
Check existing functionalities
Agree on requirements
Present proposal to Architects for approval
Define Proposal
Discuss / Review with Product and Architecture
Requires Architect approval
Approved and signed-off by Architects
Backlog grooming
Add Additional User Stories to backlog
Estimate User Story
Additional work required on story
User Story ready for sprint
Story approved and added to the backlog
Cross team planning
Sprint planning
Start work on code, test, networking, infrastructure tasks
Add story or tasks
Added to the backlog for later work
Test locally and Demo
Peer review
Requires urgent action
Check-in code
Passed review
Requires re-work
Re-factor
Raise defect
Issue detected during testing
Unit test gate
Passed unit test gate
CI gate
Passed CI gate
Package Release Build
Trigger Pipeline
Deploy to QA environment and Test
Passed QA gate
Promoted From QA
Deploy To Integration environment and test
Passed Integration gate
Promoted from Integration
Deploy To Performance environment and test
Passed Performance gate
Promoted from Test
Manual Promotion To Production
Deploy To Producton environment
Passed Production gate
Production Release Successful

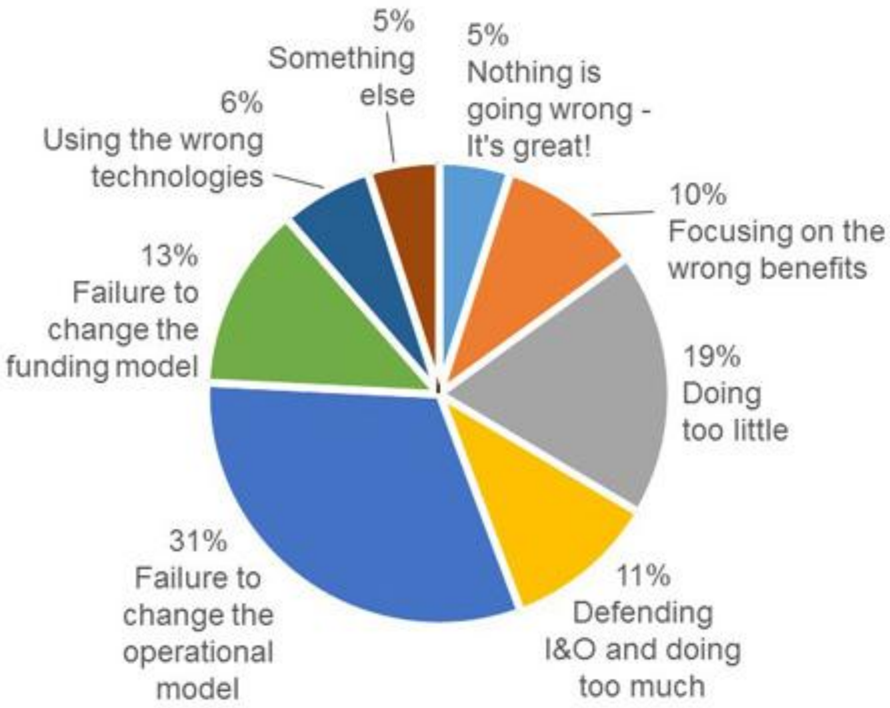
5%
Something else
5%
Nothing is going wrong - It's great!
6%
Using the wrong technologies
10%
Focusing on the wrong benefits
13%
Failure to change the funding model
19%
Doing too little
31%
Failure to change the operational model
11%
Defending I&O and doing too much

Chapter 4: Configuring Network Devices Using Ansible

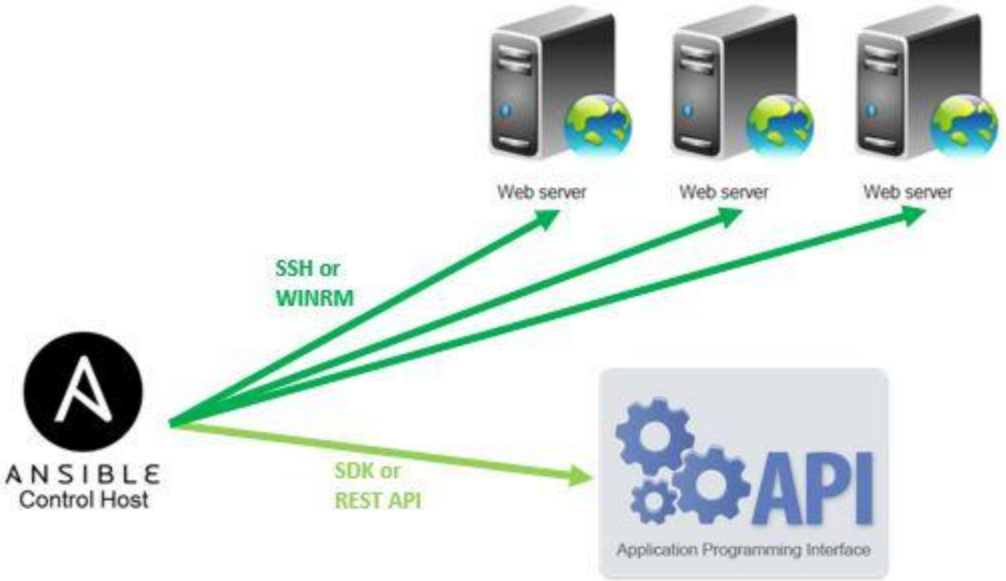

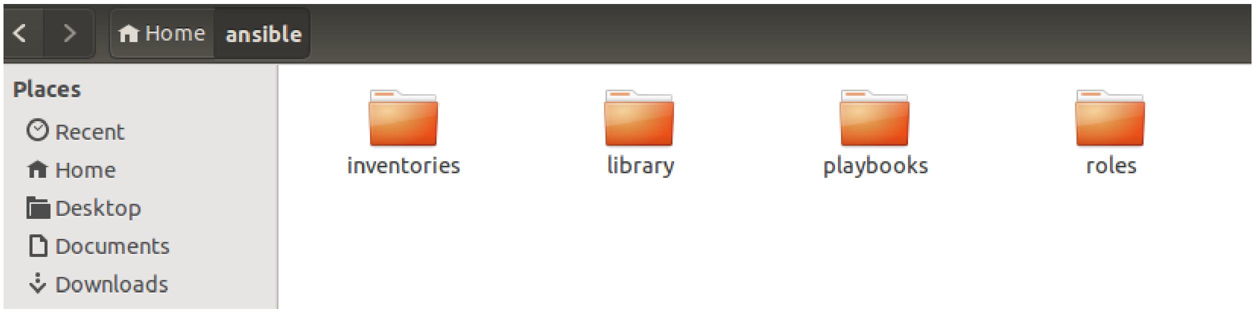

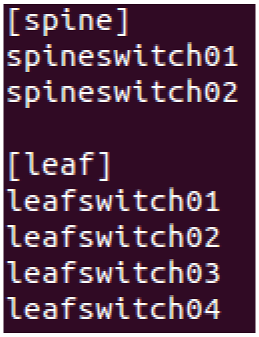

```
[spine]
spineswitch01
spineswitch02

[leaf]
leafswitch01
leafswitch02
leafswitch03
leafswitch04
```

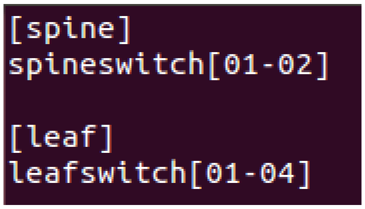

```
[spine]
spineswitch[01-02]

[leaf]
leafswitch[01-04]
```

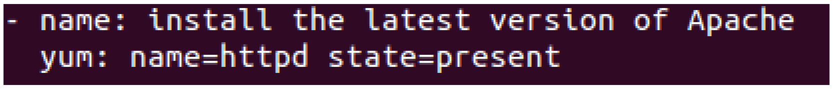

```
- name: install the latest version of Apache
  yum: name=httpd state=present
```

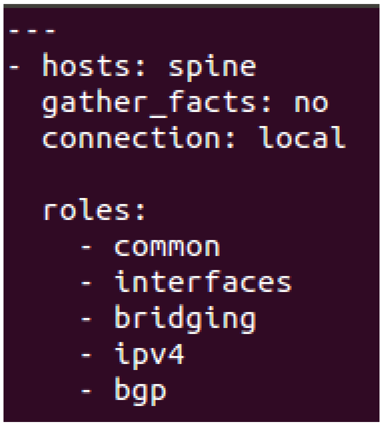

```
---
- hosts: spine
  gather_facts: no
  connection: local

  roles:
    - common
    - interfaces
    - bridging
    - ipv4
    - bgp
```

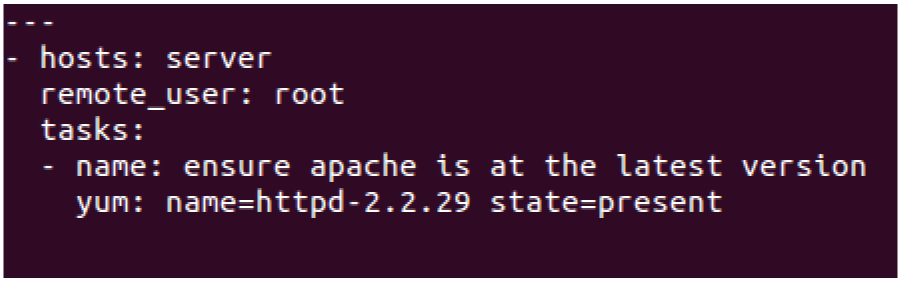

```
---
- hosts: server
  remote_user: root
  tasks:
  - name: ensure apache is at the latest version
    yum: name=httpd-2.2.29 state=present
```

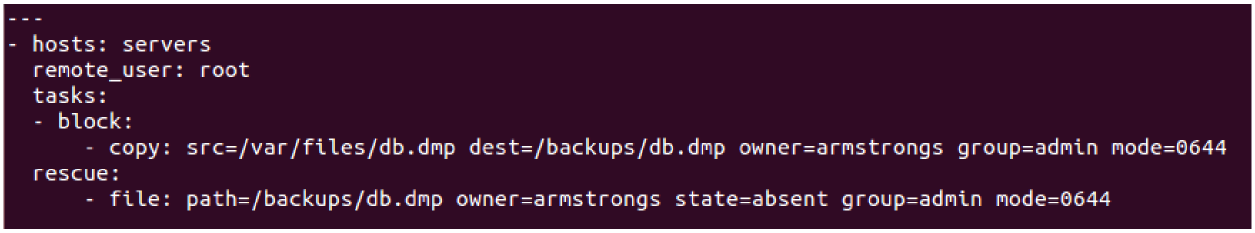

```
---
- hosts: servers
  remote_user: root
  tasks:
  - block:
      - copy: src=/var/files/db.dmp dest=/backups/db.dmp owner=armstrongs group=admin mode=0644
  rescue:
      - file: path=/backups/db.dmp owner=armstrongs state=absent group=admin mode=0644
```

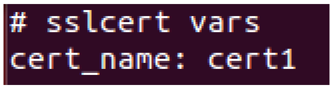

```
# sslcert vars
cert_name: cert1
```

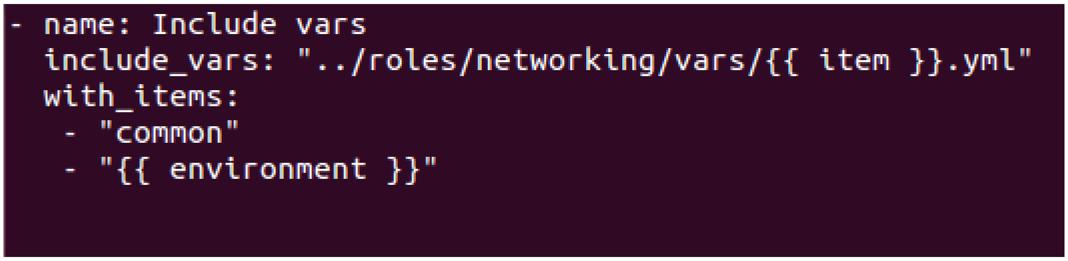

```
- name: Include vars
  include_vars: "../roles/networking/vars/{{ item }}.yml"
  with_items:
   - "common"
   - "{{ environment }}"
```

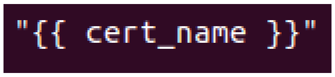

```
"{{ cert_name }}"
```

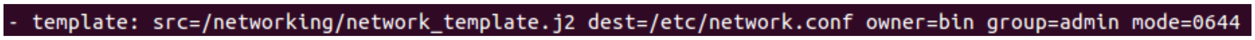

```
- template: src=/networking/network_template.j2 dest=/etc/network.conf owner=bin group=admin mode=0644
```

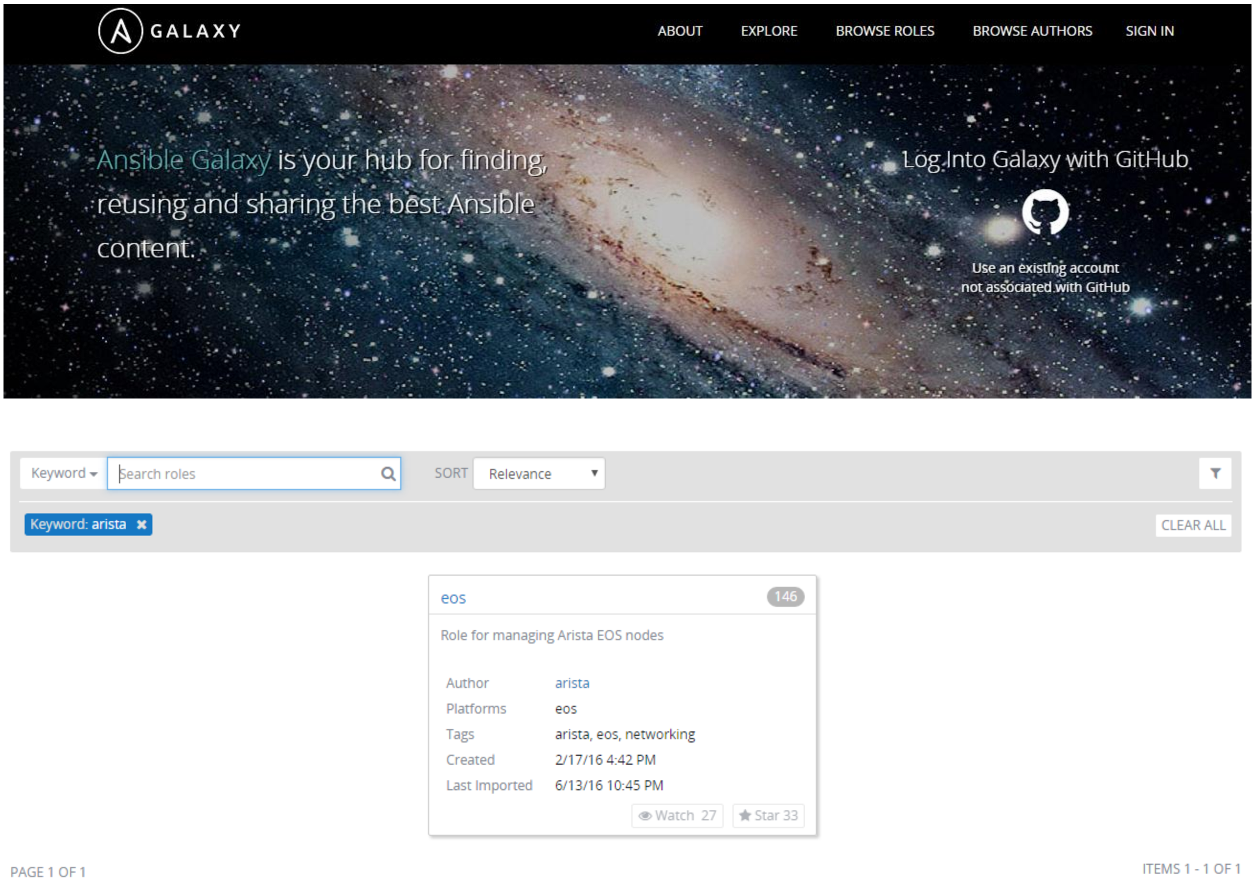

Ansible Galaxy is your hub for finding, reusing and sharing the best Ansible content.

Log Into Galaxy with GitHub

Use an existing account not associated with GitHub

eos 146

Role for managing Arista EOS nodes

Author	arista
Platforms	eos
Tags	arista, eos, networking
Created	2/17/16 4:42 PM
Last Imported	6/13/16 10:45 PM

PAGE 1 OF 1

ITEMS 1 - 1 OF 1

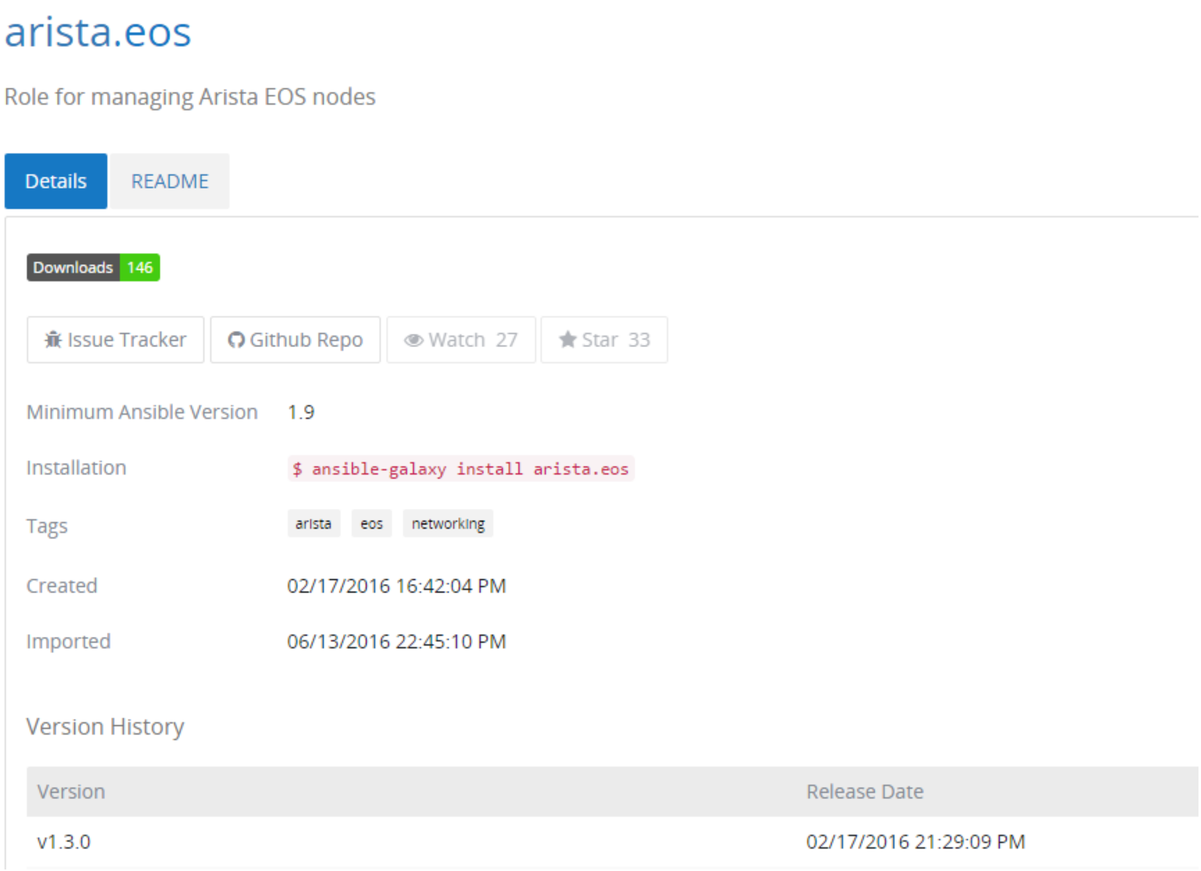

arista.eos

Role for managing Arista EOS nodes

Minimum Ansible Version 1.9

Installation `$ ansible-galaxy install arista.eos`

Tags arista eos networking

Created 02/17/2016 16:42:04 PM

Imported 06/13/2016 22:45:10 PM

Version History

Version	Release Date
v1.3.0	02/17/2016 21:29:09 PM

Junos

- junos_command - Execute arbitrary commands on a remote device running Junos
- junos_config - Manage configuration on remote devices running Junos
- junos_facts - Collect facts from remote device running Junos
- junos_netconf - Configures the Junos Netconf system service
- junos_package - Installs packages on remote devices running Junos
- junos_template - Manage configuration on remote devices running Junos

Eos

- eos_command - Run arbitrary command on EOS device
- eos_config - Manage Arista EOS configuration sections
- eos_eapi - Manage and configure EAPI. Requires EOS v4.12 or greater.
- eos_template - Manage Arista EOS device configurations

Nxos

- nxos_command - Run arbitrary command on Cisco NXOS devices
- nxos_config - Manage Cisco NXOS configuration sections
- nxos_facts - Gets facts about NX-OS switches
- nxos_feature - Manage features in NX-OS switches
- nxos_interface - Manages physical attributes of interfaces
- nxos_ip_interface - Manages L3 attributes for IPv4 and IPv6 interfaces
- nxos_nxapi - Manage NXAPI configuration on an NXOS device.
- nxos_ping - Tests reachability using ping from Nexus switch
- nxos_switchport - Manages Layer 2 switchport interfaces
- nxos_template - Manage Cisco NXOS device configurations
- nxos_vlan - Manages VLAN resources and attributes
- nxos_vrf - Manages global VRF configuration
- nxos_vrf_interface - Manages interface specific VRF configuration
- nxos_vrrp - Manages VRRP configuration on NX-OS switches

Ios

- ios_command - Run arbitrary commands on ios devices.
- ios_config - Manage Cisco IOS configuration sections
- ios_template - Manage Cisco IOS device configurations over SSH

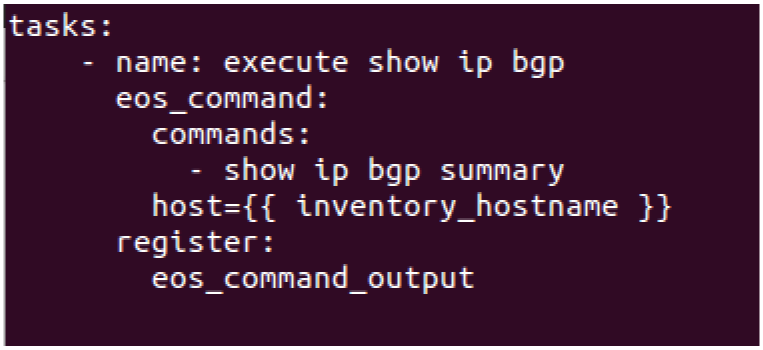

```
tasks:
    - name: execute show ip bgp
      eos_command:
        commands:
          - show ip bgp summary
        host={{ inventory_hostname }}
      register:
        eos_command_output
```

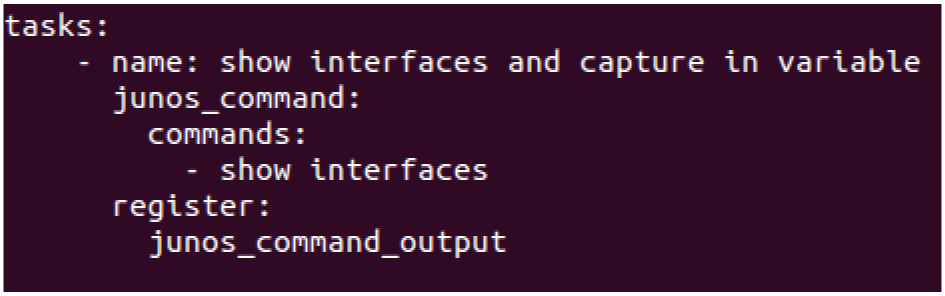

```
tasks:
    - name: show interfaces and capture in variable
      junos_command:
        commands:
          - show interfaces
      register:
        junos_command_output
```

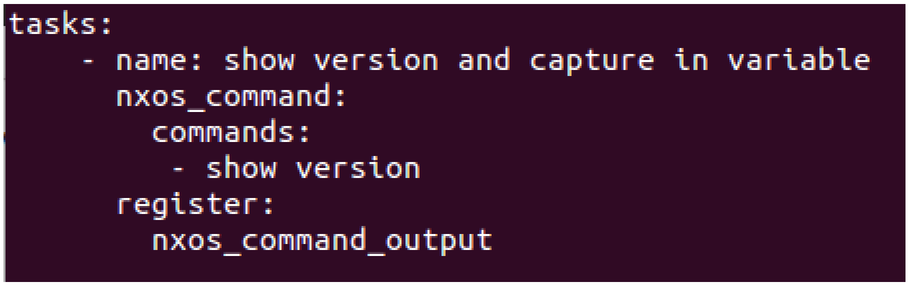

```
tasks:
    - name: show version and capture in variable
      nxos_command:
        commands:
         - show version
      register:
        nxos_command_output
```

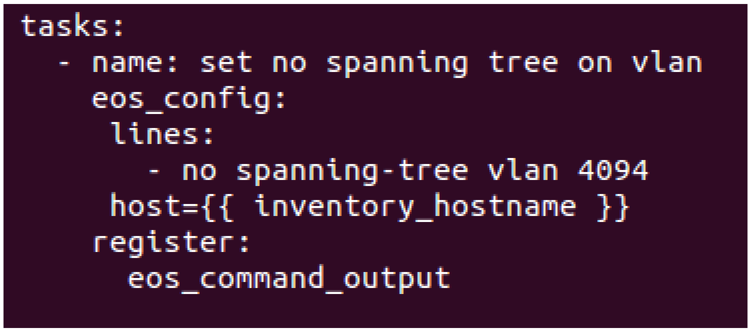

```
tasks:
  - name: set no spanning tree on vlan
    eos_config:
     lines:
       - no spanning-tree vlan 4094
     host={{ inventory_hostname }}
    register:
      eos_command_output
```

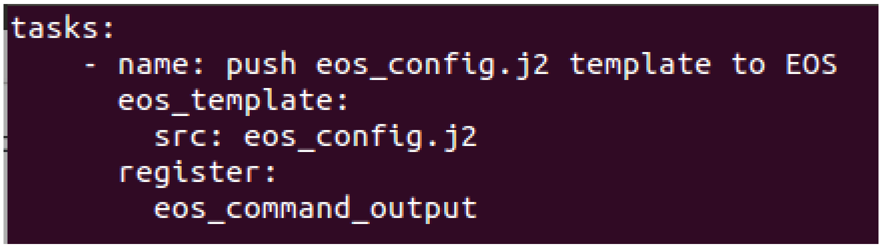

```
tasks:
    - name: push eos_config.j2 template to EOS
      eos_template:
        src: eos_config.j2
      register:
        eos_command_output
```

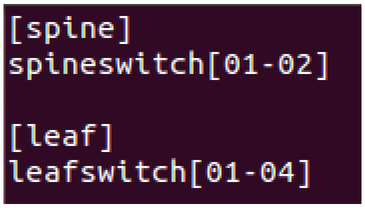

```
[spine]
spineswitch[01-02]

[leaf]
leafswitch[01-04]
```

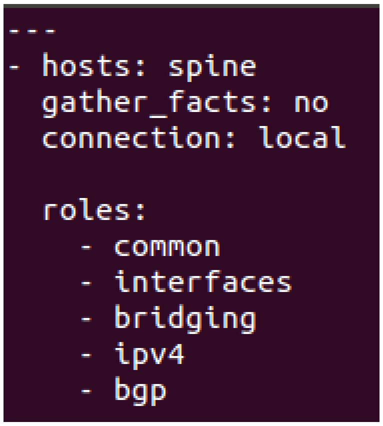

```
---
- hosts: spine
  gather_facts: no
  connection: local

  roles:
    - common
    - interfaces
    - bridging
    - ipv4
    - bgp
```

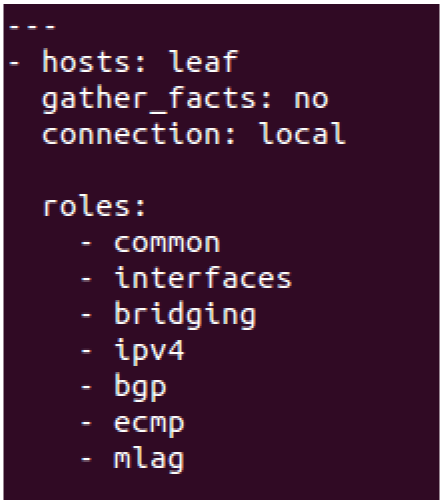

```
---
- hosts: leaf
  gather_facts: no
  connection: local

  roles:
    - common
    - interfaces
    - bridging
    - ipv4
    - bgp
    - ecmp
    - mlag
```

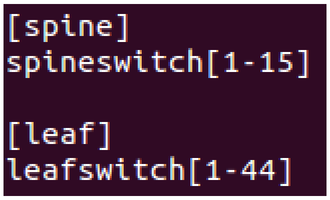

```
[spine]
spineswitch[1-15]

[leaf]
leafswitch[1-44]
```

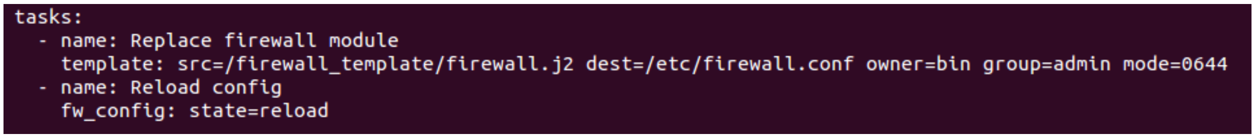

```
tasks:
  - name: Replace firewall module
    template: src=/firewall_template/firewall.j2 dest=/etc/firewall.conf owner=bin group=admin mode=0644
  - name: Reload config
    fw_config: state=reload
```

Chapter 5: Orchestrating Load Balancers Using Ansible

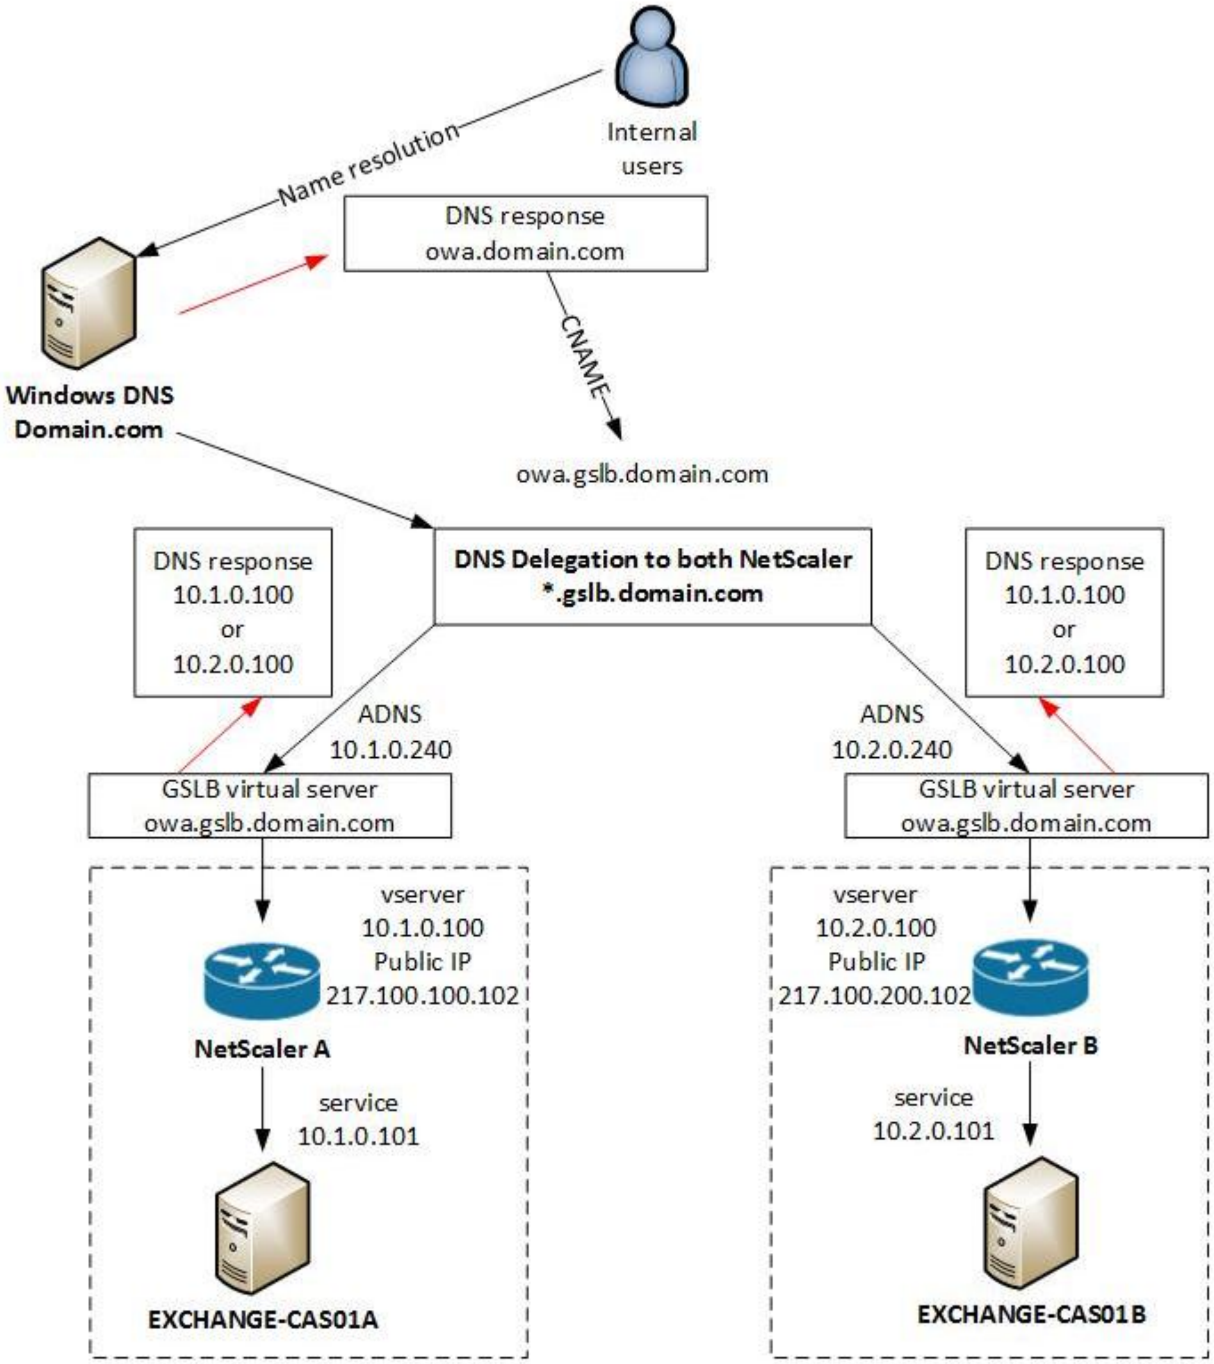

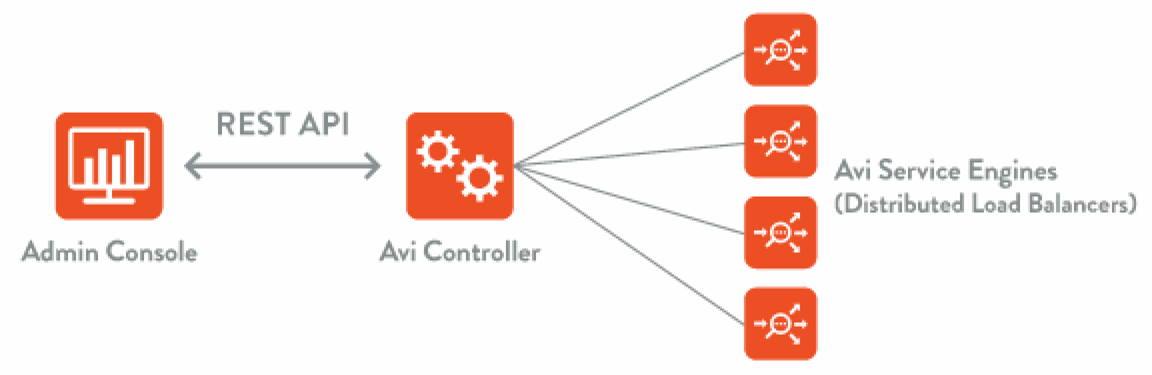

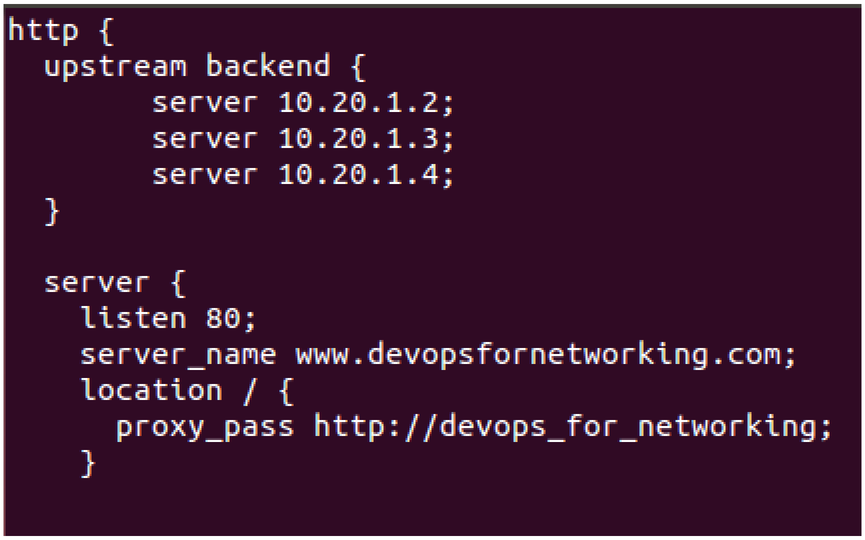

```
http {
  upstream backend {
        server 10.20.1.2;
        server 10.20.1.3;
        server 10.20.1.4;
  }

  server {
    listen 80;
    server_name www.devopsfornetworking.com;
    location / {
      proxy_pass http://devops_for_networking;
    }
```

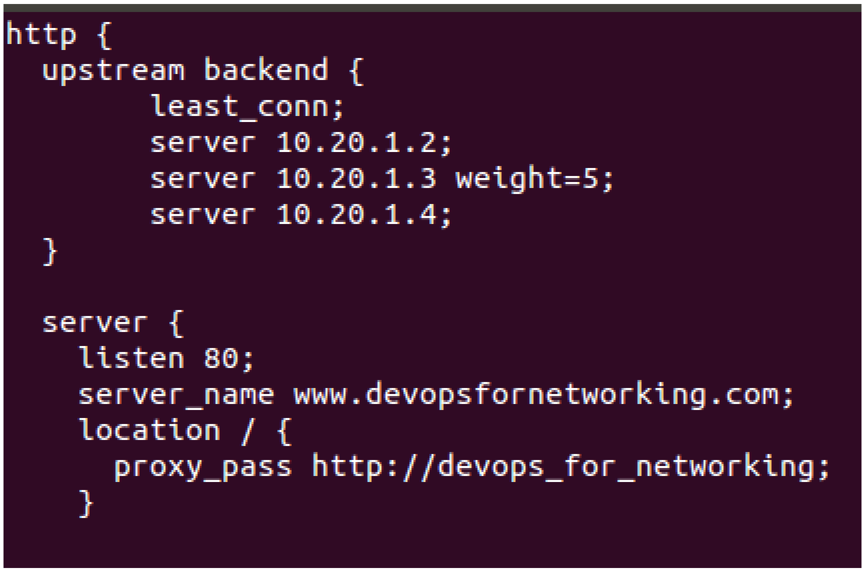

```
http {
  upstream backend {
        least_conn;
        server 10.20.1.2;
        server 10.20.1.3 weight=5;
        server 10.20.1.4;
  }

  server {
    listen 80;
    server_name www.devopsfornetworking.com;
    location / {
      proxy_pass http://devops_for_networking;
    }
```

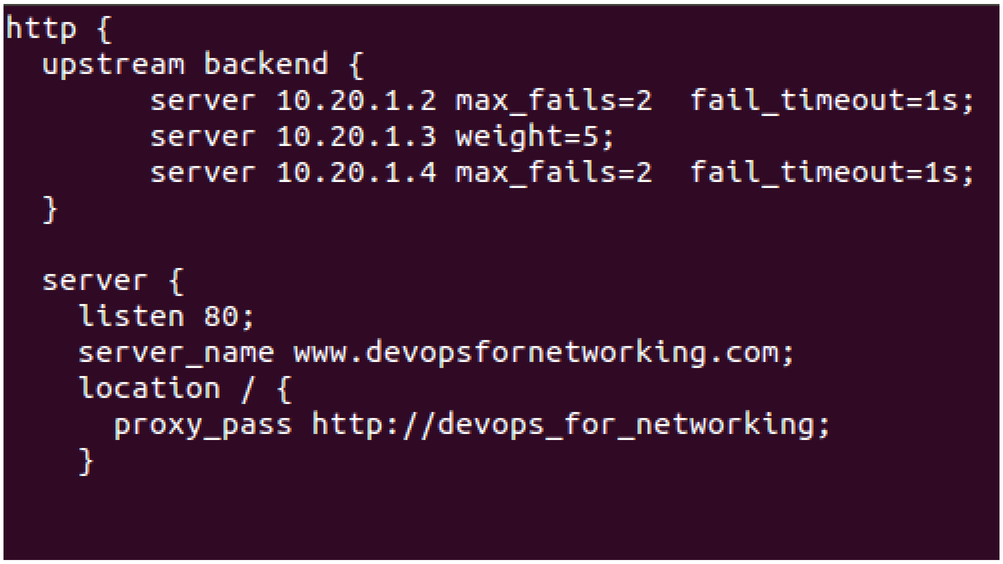

```
http {
  upstream backend {
        server 10.20.1.2 max_fails=2  fail_timeout=1s;
        server 10.20.1.3 weight=5;
        server 10.20.1.4 max_fails=2  fail_timeout=1s;
  }

  server {
    listen 80;
    server_name www.devopsfornetworking.com;
    location / {
      proxy_pass http://devops_for_networking;
    }
```

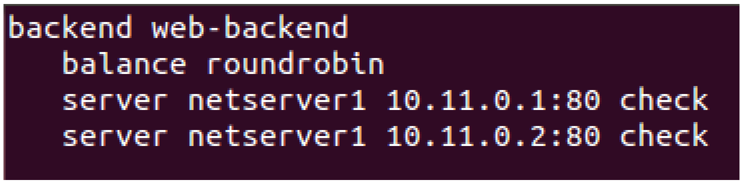

```
backend web-backend
   balance roundrobin
   server netserver1 10.11.0.1:80 check
   server netserver1 10.11.0.2:80 check
```

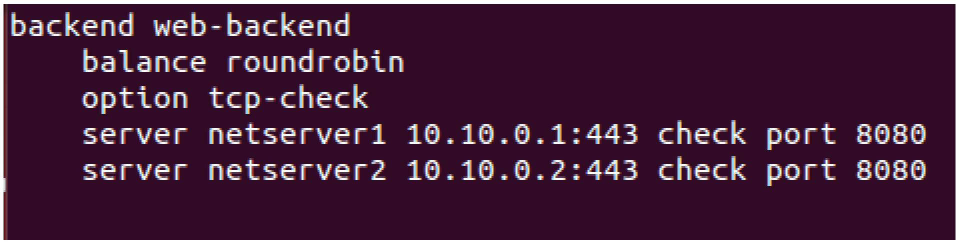

```
backend web-backend
    balance roundrobin
    option tcp-check
    server netserver1 10.10.0.1:443 check port 8080
    server netserver2 10.10.0.2:443 check port 8080
```

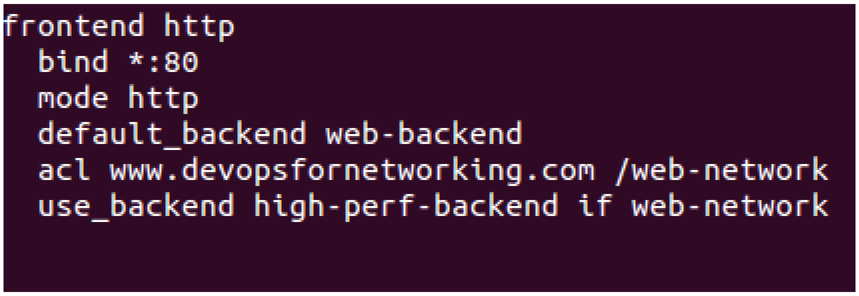

```
frontend http
  bind *:80
  mode http
  default_backend web-backend
  acl www.devopsfornetworking.com /web-network
  use_backend high-perf-backend if web-network
```

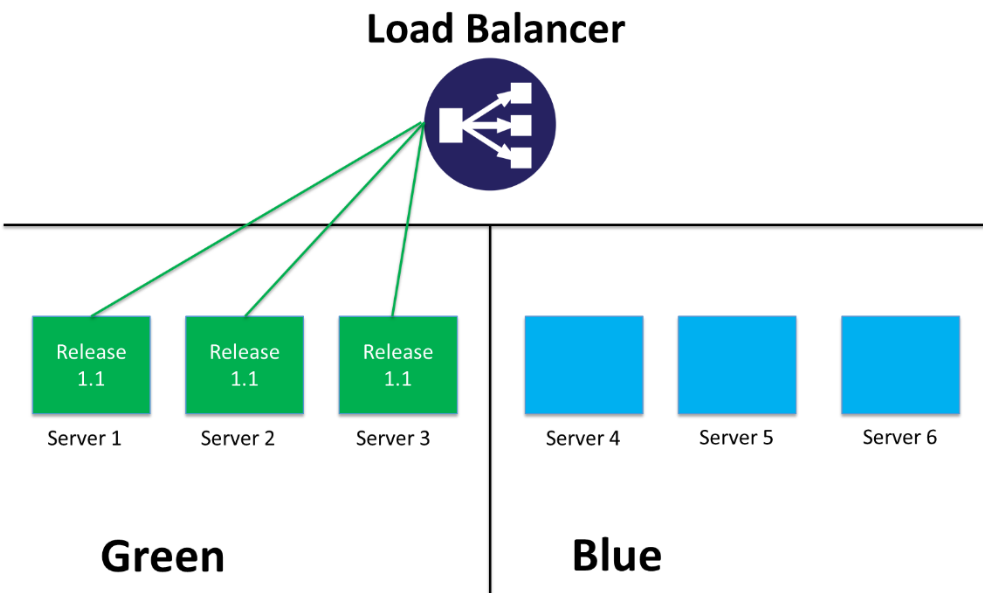
Load Balancer
Release 1.1
Release 1.1
Release 1.1
Server 1
Server 2
Server 3
Server 4
Server 5
Server 6
Green
Blue

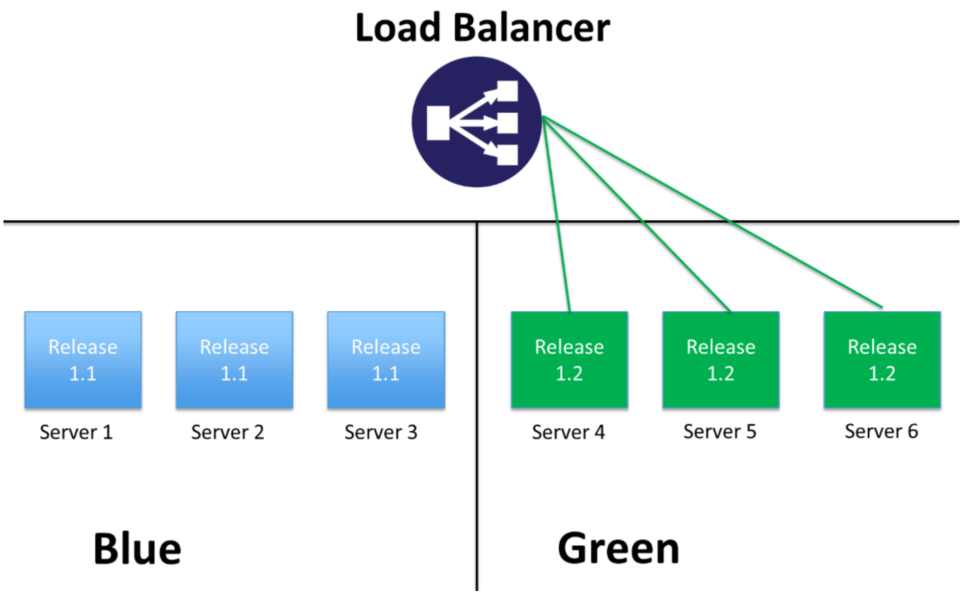
Load Balancer
Release 1.1
Release 1.1
Release 1.1
Release 1.2
Release 1.2
Release 1.2
Server 1
Server 2
Server 3
Server 4
Server 5
Server 6
Blue
Green

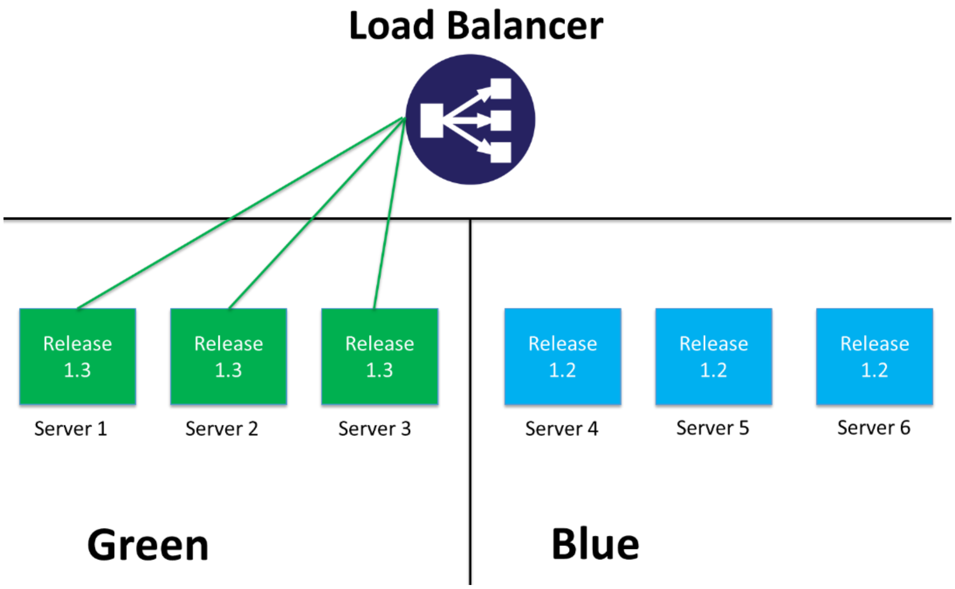
Load Balancer
Release 1.3
Server 1
Release 1.3
Server 2
Release 1.3
Server 3
Release 1.2
Server 4
Release 1.2
Server 5
Release 1.2
Server 6
Green
Blue

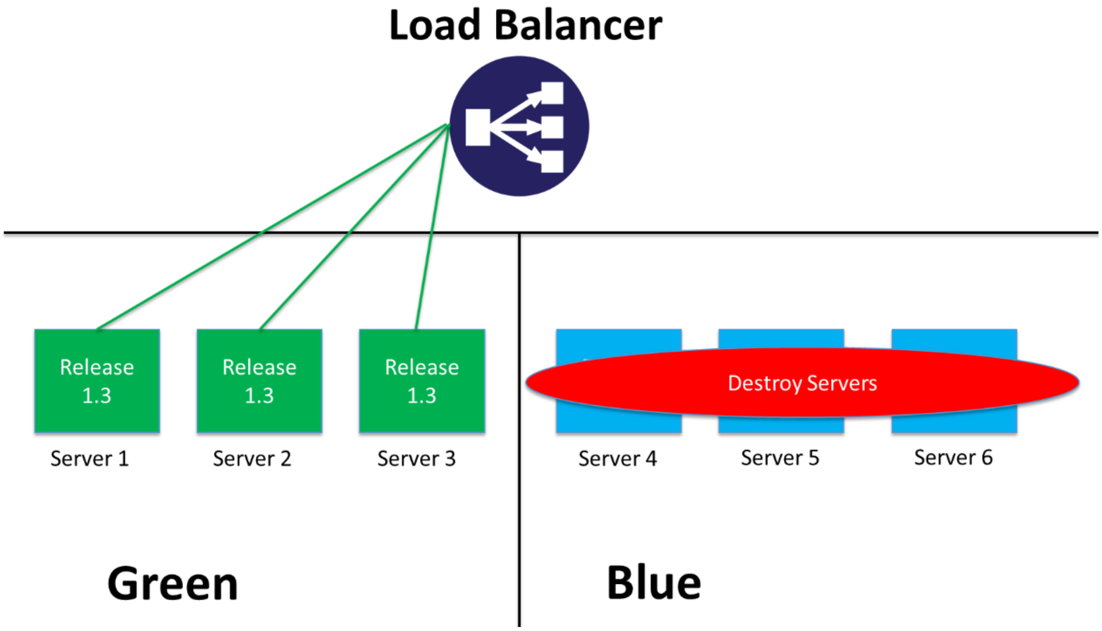
Load Balancer
Release 1.3
Server 1
Release 1.3
Server 2
Release 1.3
Server 3
Destroy Servers
Server 4
Server 5
Server 6
Green
Blue

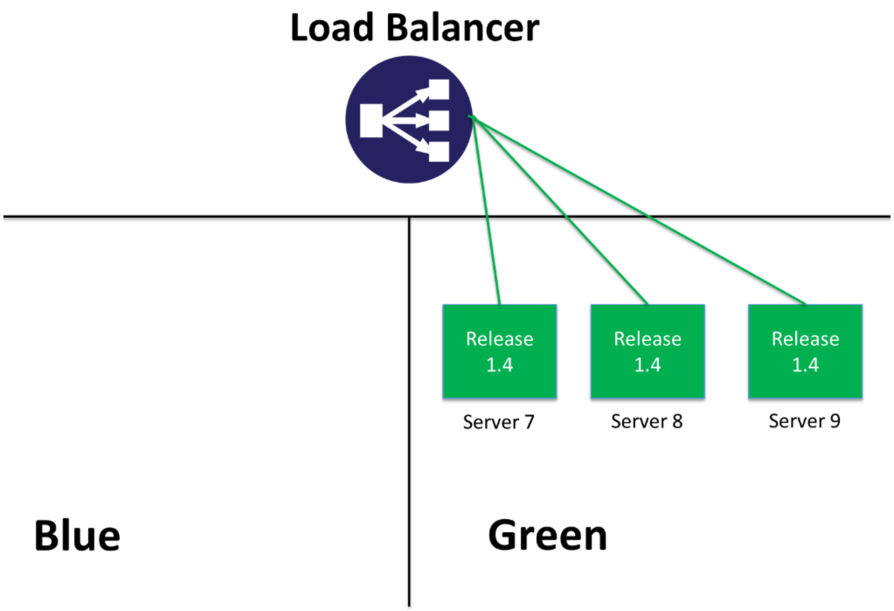

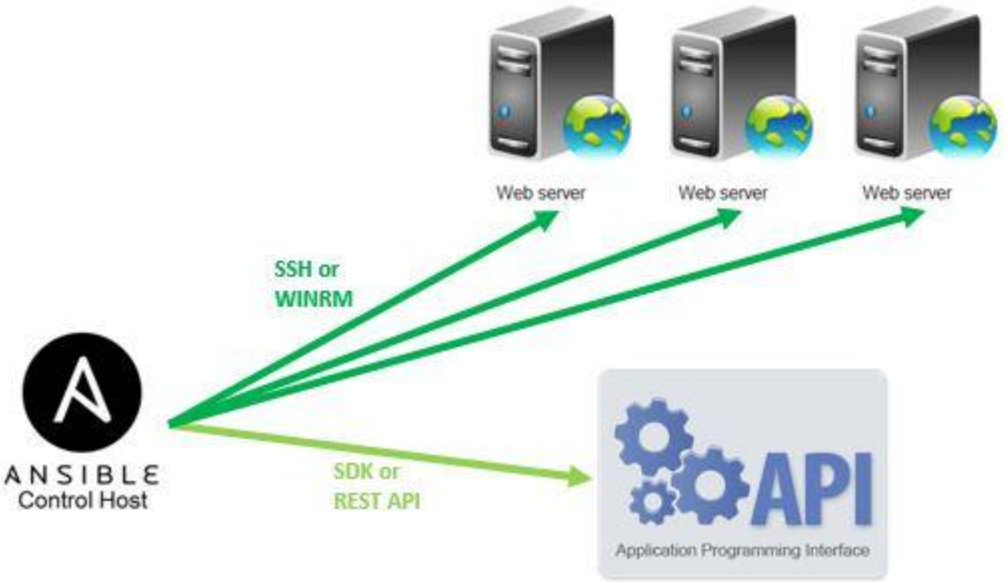

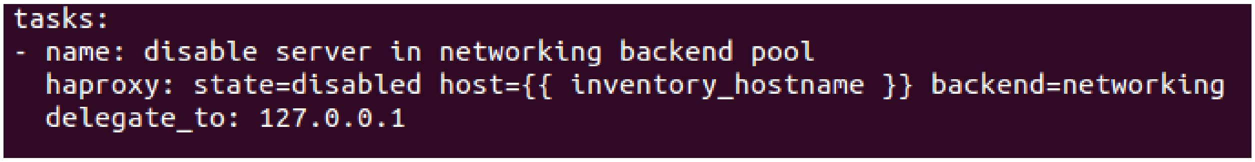

```
tasks:
- name: disable server in networking backend pool
  haproxy: state=disabled host={{ inventory_hostname }} backend=networking
  delegate_to: 127.0.0.1
```

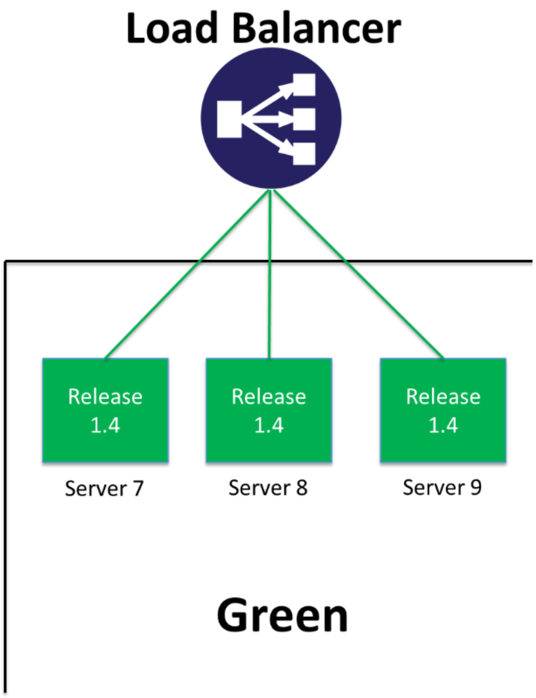

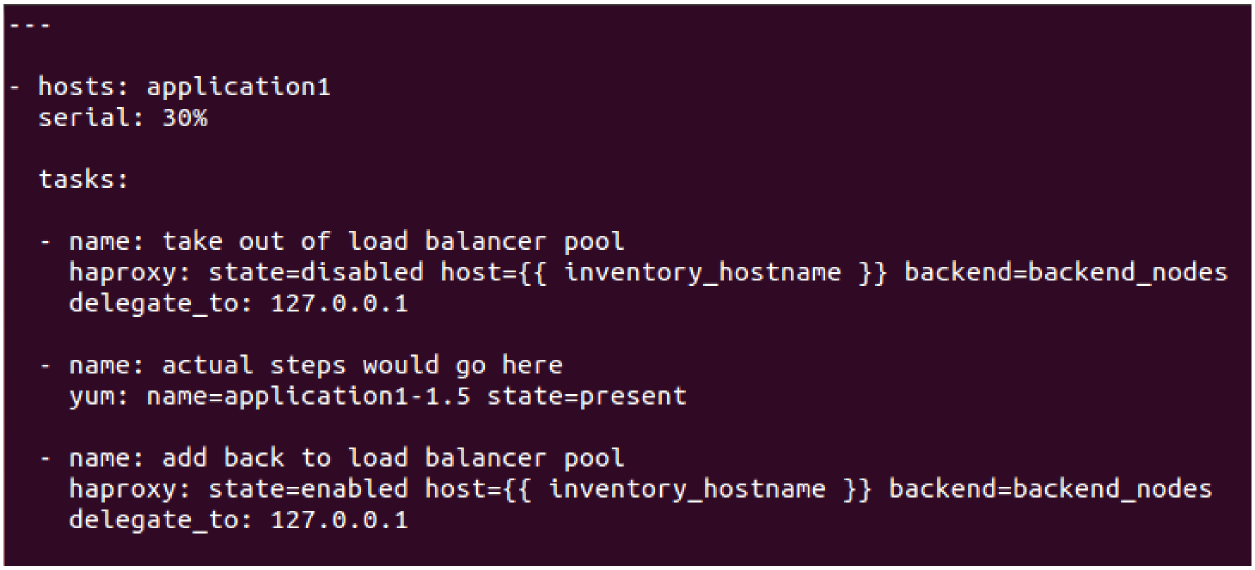

```
---

- hosts: application1
  serial: 30%

  tasks:

  - name: take out of load balancer pool
    haproxy: state=disabled host={{ inventory_hostname }} backend=backend_nodes
    delegate_to: 127.0.0.1

  - name: actual steps would go here
    yum: name=application1-1.5 state=present

  - name: add back to load balancer pool
    haproxy: state=enabled host={{ inventory_hostname }} backend=backend_nodes
    delegate_to: 127.0.0.1
```

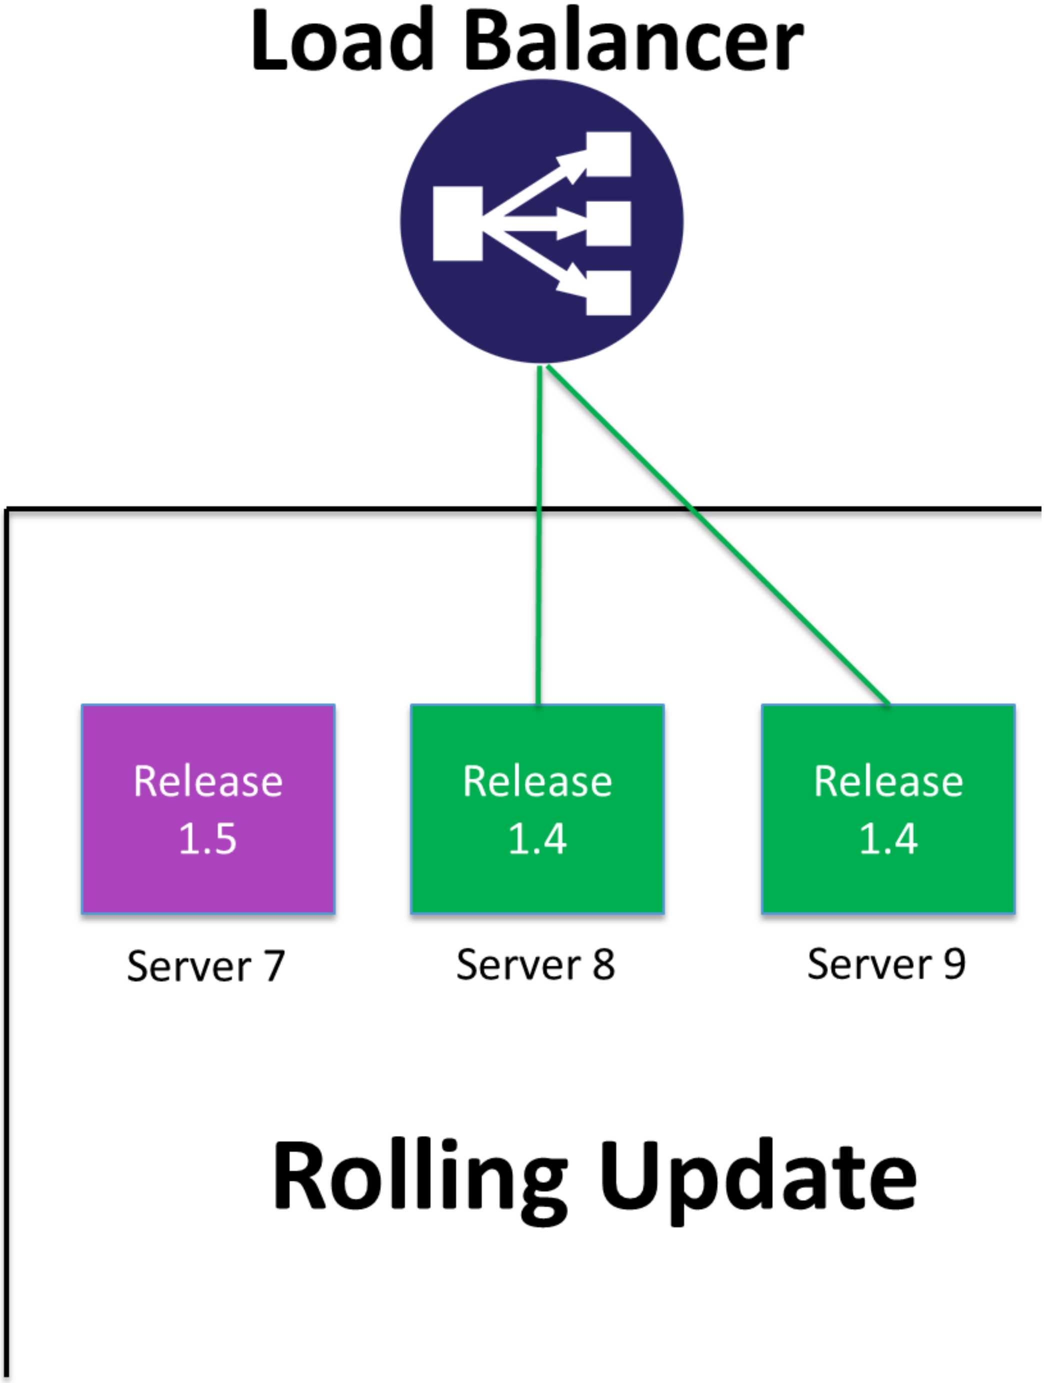

Load Balancer

Release 1.5

Server 7

Release 1.4

Server 8

Release 1.4

Server 9

Rolling Update

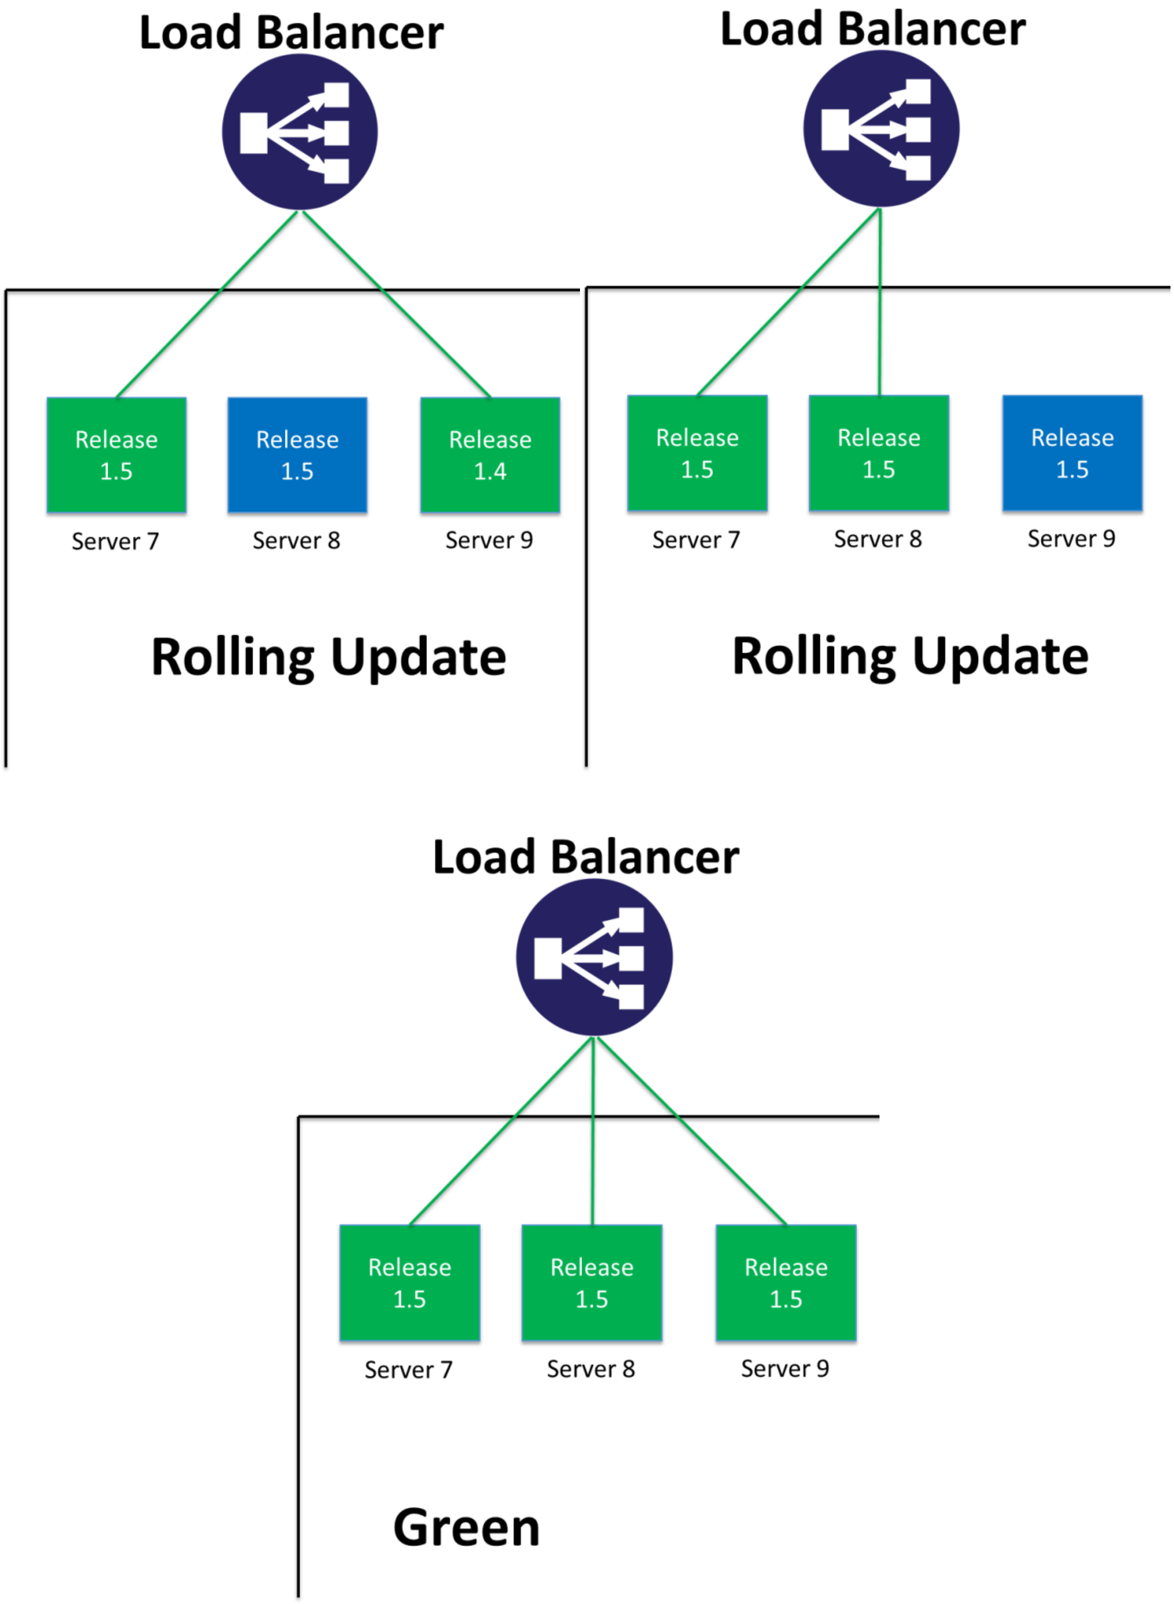
Load Balancer
Release 1.5
Server 7
Release 1.5
Server 8
Release 1.4
Server 9
Rolling Update
Load Balancer
Release 1.5
Server 7
Release 1.5
Server 8
Release 1.5
Server 9
Rolling Update
Load Balancer
Release 1.5
Server 7
Release 1.5
Server 8
Release 1.5
Server 9
Green

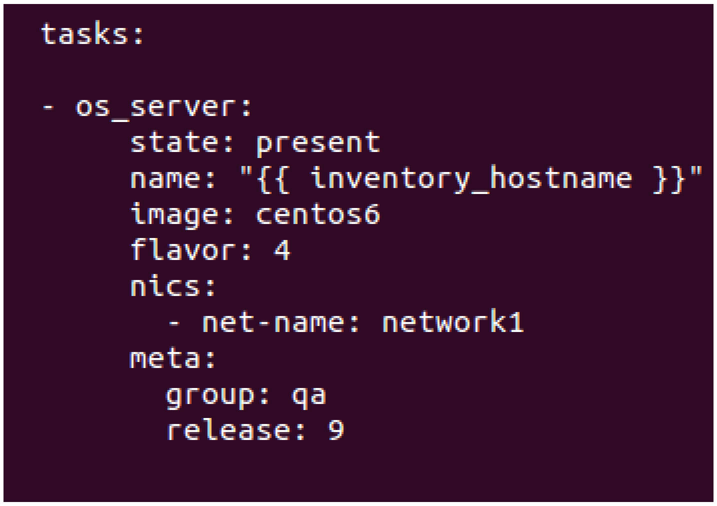

```
tasks:

- os_server:
     state: present
     name: "{{ inventory_hostname }}"
     image: centos6
     flavor: 4
     nics:
       - net-name: network1
     meta:
       group: qa
       release: 9
```

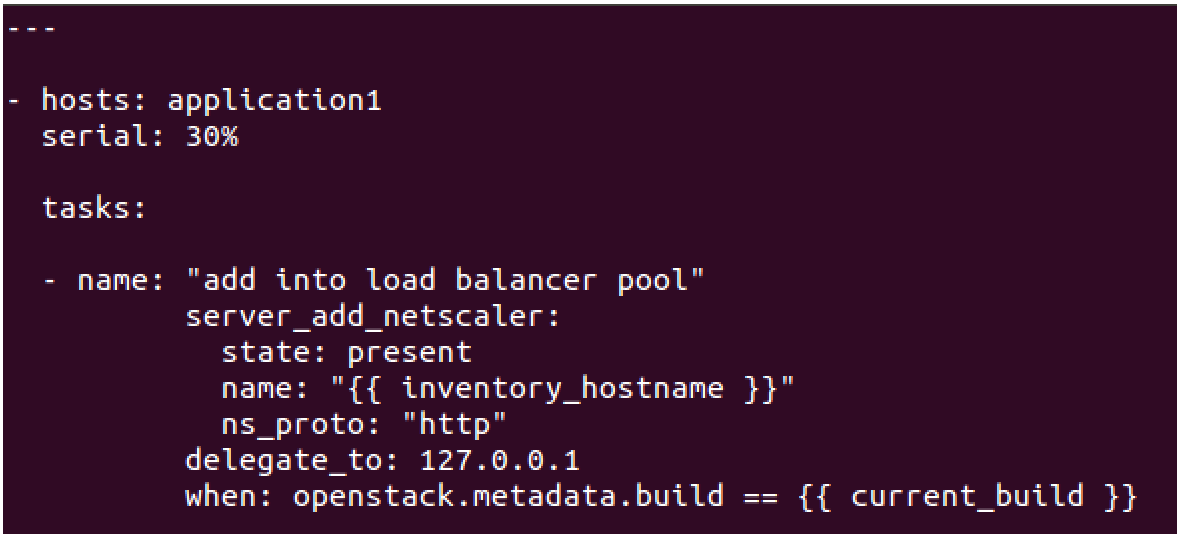

```
---

- hosts: application1
  serial: 30%

  tasks:

  - name: "add into load balancer pool"
          server_add_netscaler:
            state: present
            name: "{{ inventory_hostname }}"
            ns_proto: "http"
          delegate_to: 127.0.0.1
          when: openstack.metadata.build == {{ current_build }}
```

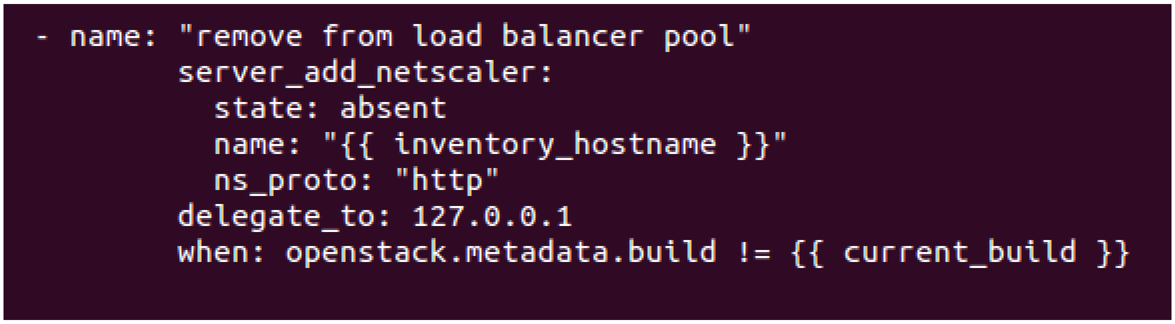

```
  - name: "remove from load balancer pool"
          server_add_netscaler:
            state: absent
            name: "{{ inventory_hostname }}"
            ns_proto: "http"
          delegate_to: 127.0.0.1
          when: openstack.metadata.build != {{ current_build }}
```

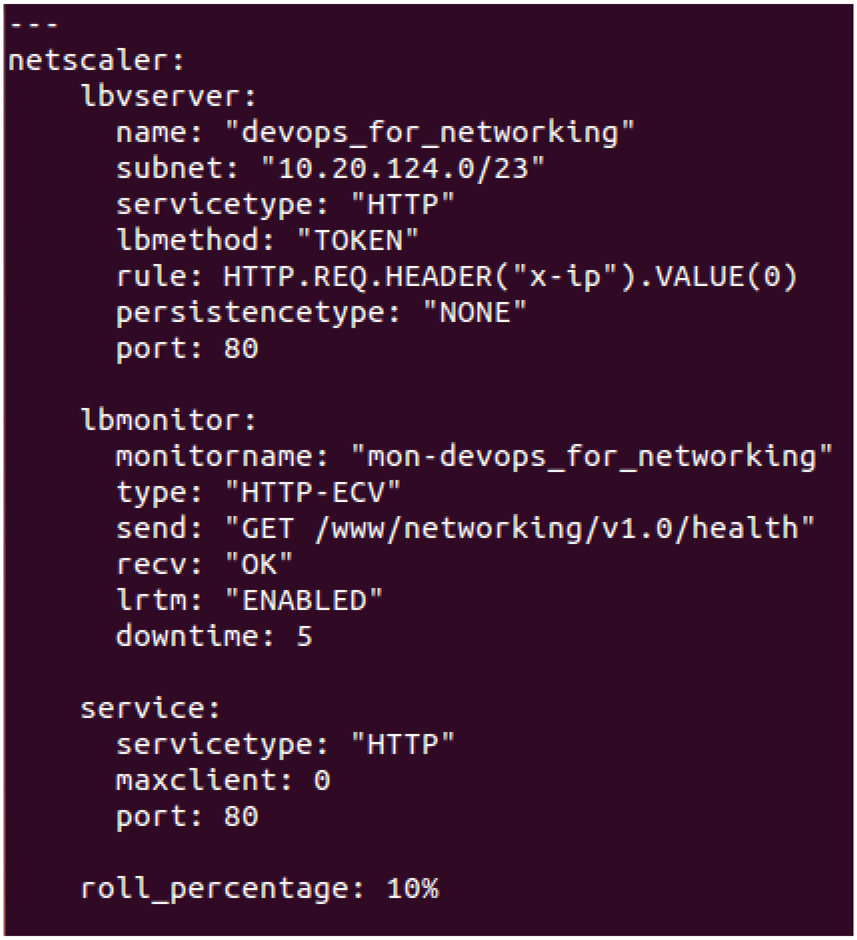

```
---
netscaler:
    lbvserver:
      name: "devops_for_networking"
      subnet: "10.20.124.0/23"
      servicetype: "HTTP"
      lbmethod: "TOKEN"
      rule: HTTP.REQ.HEADER("x-ip").VALUE(0)
      persistencetype: "NONE"
      port: 80

    lbmonitor:
      monitorname: "mon-devops_for_networking"
      type: "HTTP-ECV"
      send: "GET /www/networking/v1.0/health"
      recv: "OK"
      lrtm: "ENABLED"
      downtime: 5

    service:
      servicetype: "HTTP"
      maxclient: 0
      port: 80

    roll_percentage: 10%
```

Chapter 6: Orchestrating SDN Controllers Using Ansible

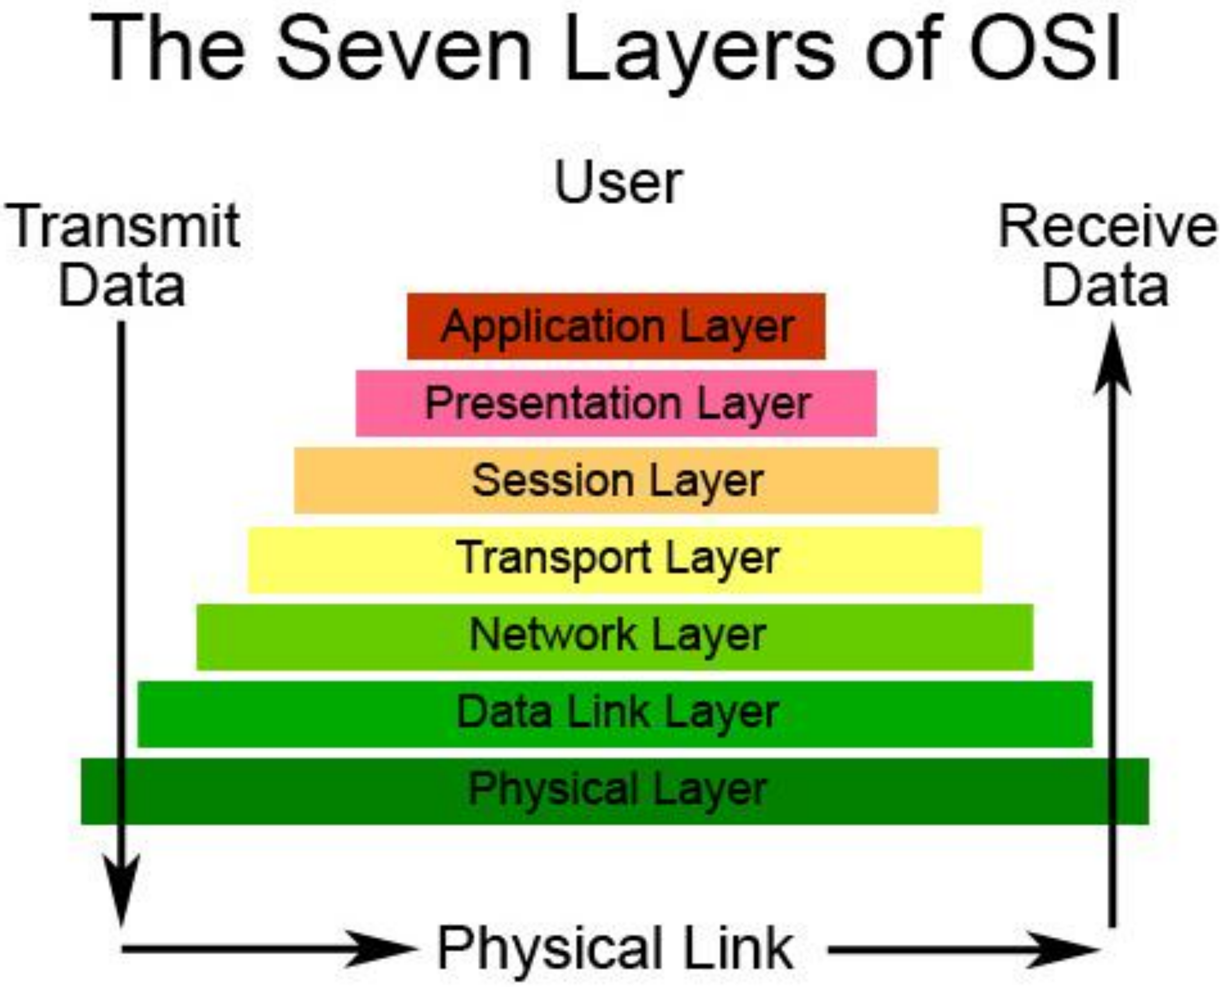

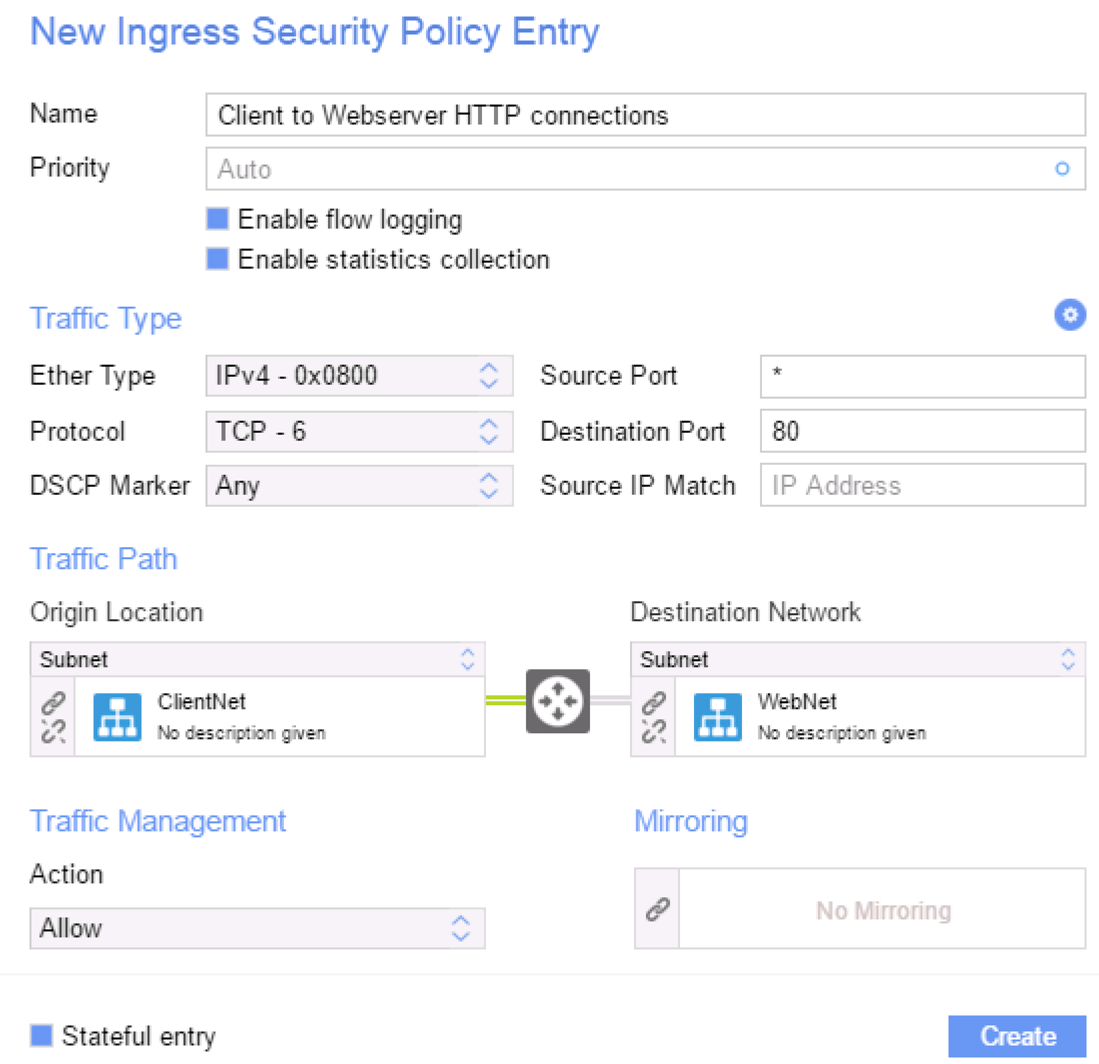

New Ingress Security Policy Entry

Name	Client to Webserver HTTP connections
Priority	Auto

■ Enable flow logging

■ Enable statistics collection

Traffic Type

Ether Type	IPv4 - 0x0800	Source Port	*
Protocol	TCP - 6	Destination Port	80
DSCP Marker	Any	Source IP Match	IP Address

Traffic Path

Origin Location: Subnet — ClientNet, No description given

Destination Network: Subnet — WebNet, No description given

Traffic Management

Action: Allow

Mirroring

No Mirroring

■ Stateful entry

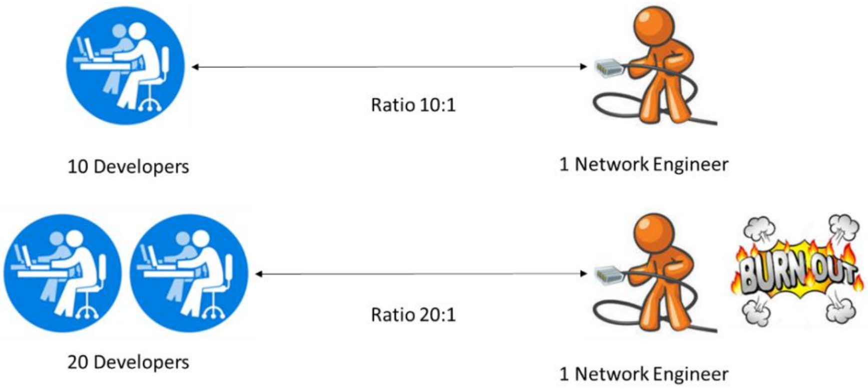

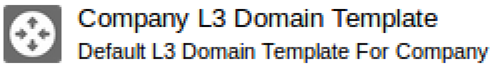

Company L3 Domain Template
Default L3 Domain Template For Company

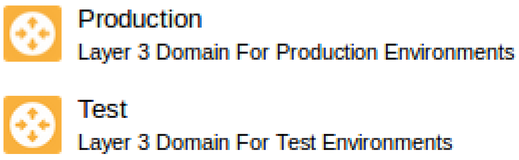

Production
Layer 3 Domain For Production Environments

Test
Layer 3 Domain For Test Environments

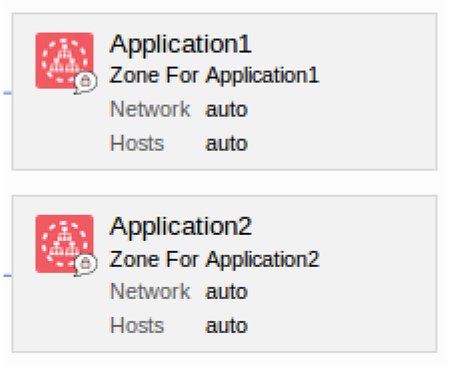

Application1
Zone For Application1
Network auto
Hosts auto

Application2
Zone For Application2
Network auto
Hosts auto

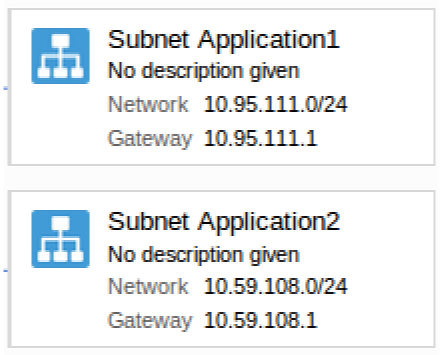

Subnet Application1
No description given
Network 10.95.111.0/24
Gateway 10.95.111.1

Subnet Application2
No description given
Network 10.59.108.0/24
Gateway 10.59.108.1

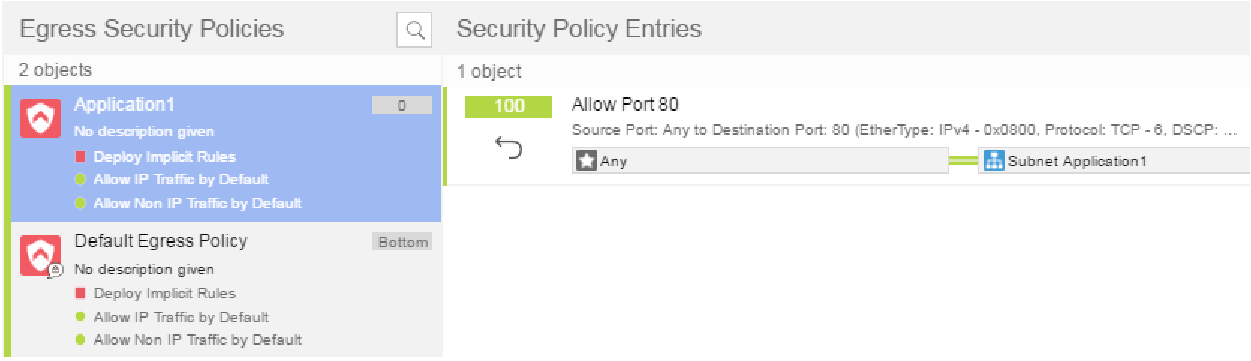

Egress Security Policies

2 objects

Application1 0
No description given
- 🟥 Deploy Implicit Rules
- ● Allow IP Traffic by Default
- ● Allow Non IP Traffic by Default

Default Egress Policy Bottom
No description given
- 🟥 Deploy Implicit Rules
- ● Allow IP Traffic by Default
- ● Allow Non IP Traffic by Default

Security Policy Entries

1 object

100 Allow Port 80
Source Port: Any to Destination Port: 80 (EtherType: IPv4 - 0x0800, Protocol: TCP - 6, DSCP: ...
Any — Subnet Application1

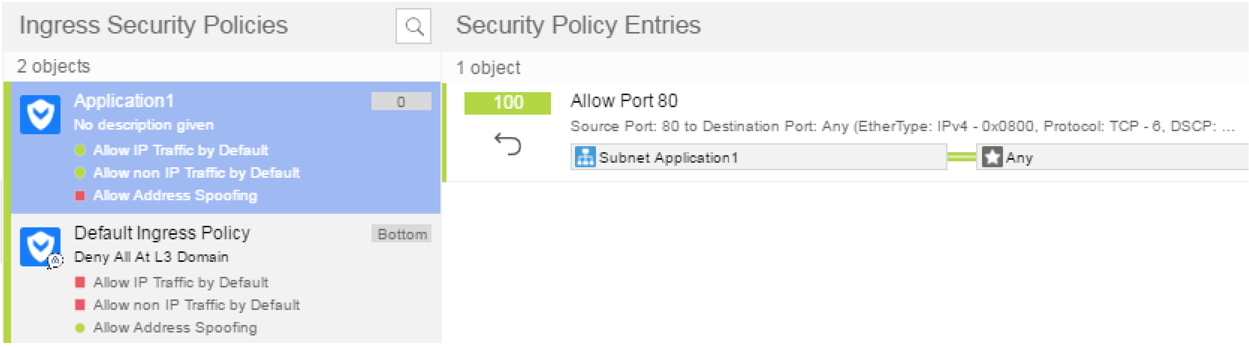

Ingress Security Policies

2 objects

Application1 — 0
No description given
● Allow IP Traffic by Default
● Allow non IP Traffic by Default
■ Allow Address Spoofing

Default Ingress Policy — Bottom
Deny All At L3 Domain
■ Allow IP Traffic by Default
■ Allow non IP Traffic by Default
● Allow Address Spoofing

Security Policy Entries

1 object

100 — Allow Port 80
Source Port: 80 to Destination Port: Any (EtherType: IPv4 - 0x0800, Protocol: TCP - 6, DSCP: ...
Subnet Application1 — Any

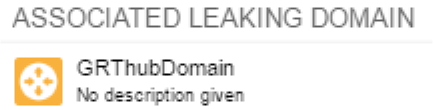

ASSOCIATED LEAKING DOMAIN

GRThubDomain
No description given

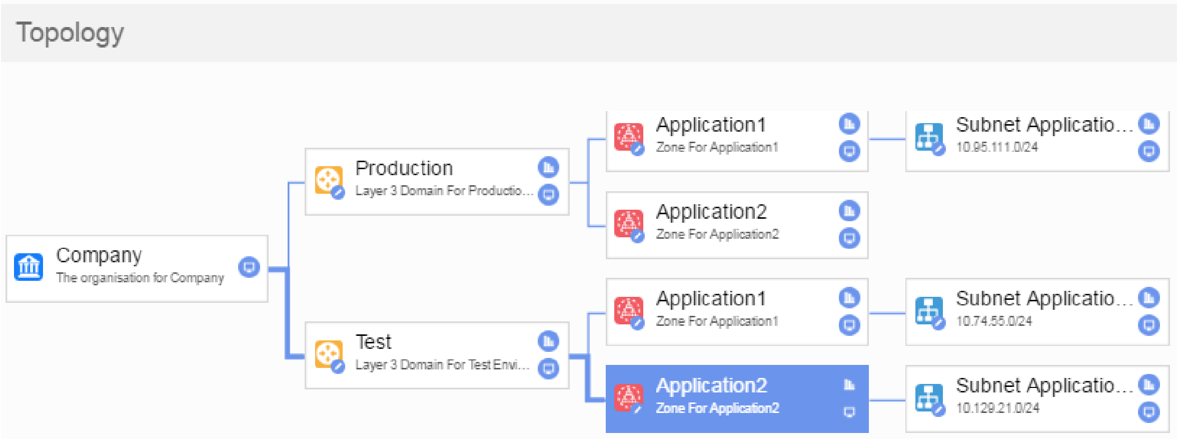

Topology

Company
The organisation for Company

Production
Layer 3 Domain For Productio...

Test
Layer 3 Domain For Test Envi...

Application1
Zone For Application1

Application2
Zone For Application2

Application1
Zone For Application1

Application2
Zone For Application2

Subnet Applicatio...
10.95.111.0/24

Subnet Applicatio...
10.74.55.0/24

Subnet Applicatio...
10.129.21.0/24

Company
The organisation for Company

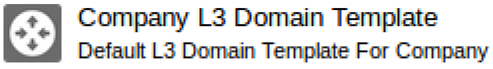

Company L3 Domain Template
Default L3 Domain Template For Company

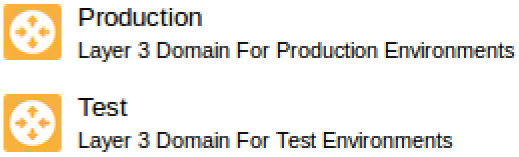

Production
Layer 3 Domain For Production Environments

Test
Layer 3 Domain For Test Environments

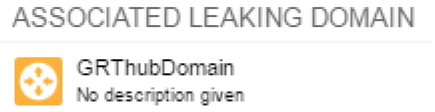

ASSOCIATED LEAKING DOMAIN

GRThubDomain
No description given

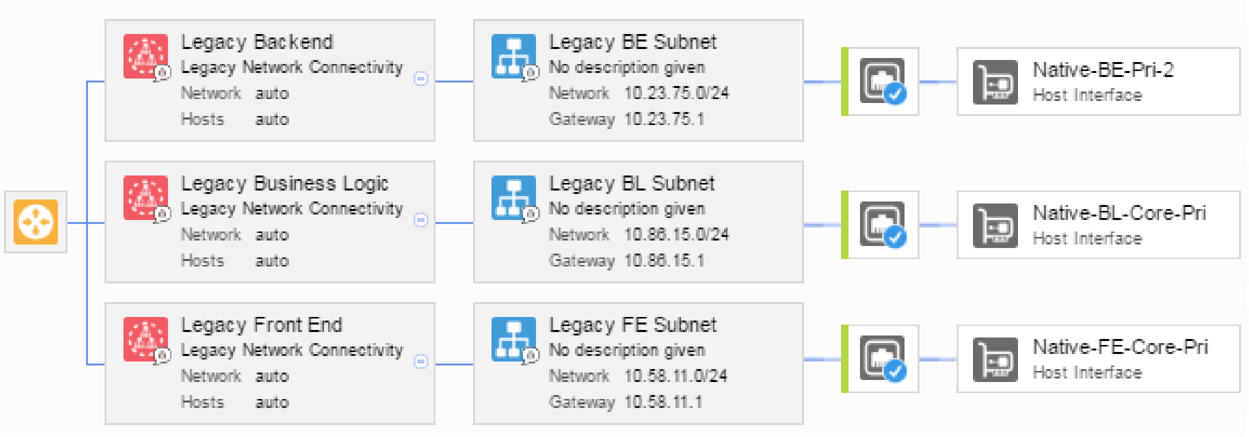

Legacy Backend
Legacy Network Connectivity
Network auto
Hosts auto

Legacy BE Subnet
No description given
Network 10.23.75.0/24
Gateway 10.23.75.1

Native-BE-Pri-2
Host Interface

Legacy Business Logic
Legacy Network Connectivity
Network auto
Hosts auto

Legacy BL Subnet
No description given
Network 10.86.15.0/24
Gateway 10.86.15.1

Native-BL-Core-Pri
Host Interface

Legacy Front End
Legacy Network Connectivity
Network auto
Hosts auto

Legacy FE Subnet
No description given
Network 10.58.11.0/24
Gateway 10.58.11.1

Native-FE-Core-Pri
Host Interface

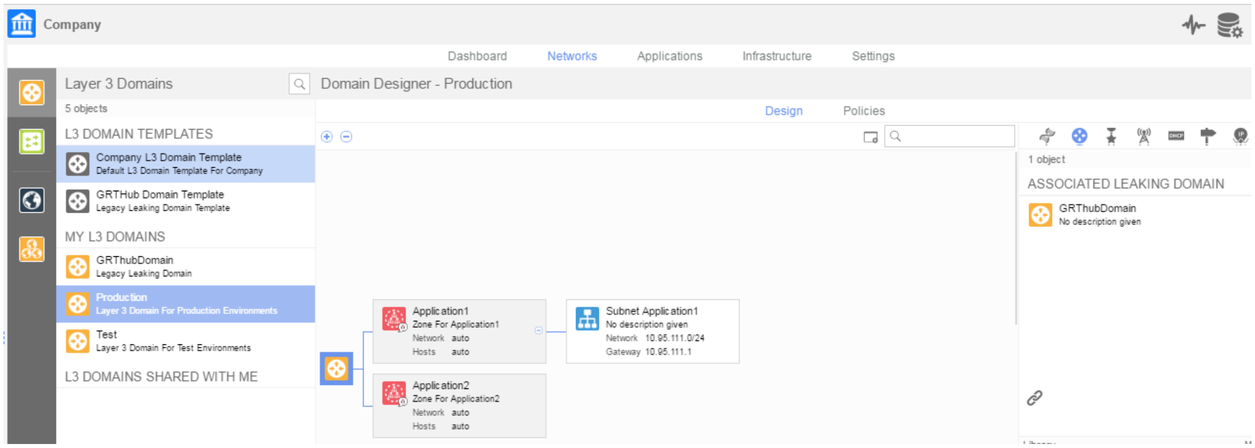

Layer 3 Domains
5 objects

L3 DOMAIN TEMPLATES

Company L3 Domain Template
Default L3 Domain Template For Company

GRTHub Domain Template
Legacy Leaking Domain Template

MY L3 DOMAINS

GRThubDomain
Legacy Leaking Domain

Production
Layer 3 Domain For Production Environments

Test
Layer 3 Domain For Test Environments

L3 DOMAINS SHARED WITH ME

Domain Designer - Production

Application1
Zone For Application1
Network auto
Hosts auto

Subnet Application1
No description given
Network 10.95.111.0/24
Gateway 10.95.111.1

Application2
Zone For Application2
Network auto
Hosts auto

1 object
ASSOCIATED LEAKING DOMAIN

GRThubDomain
No description given

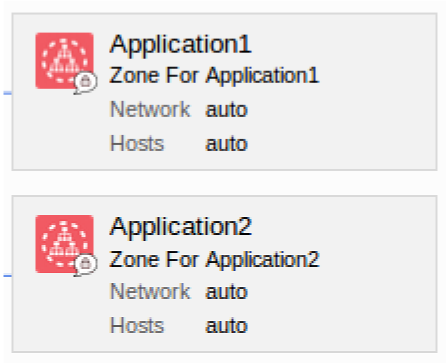

Application1
Zone For Application1
Network auto
Hosts auto

Application2
Zone For Application2
Network auto
Hosts auto

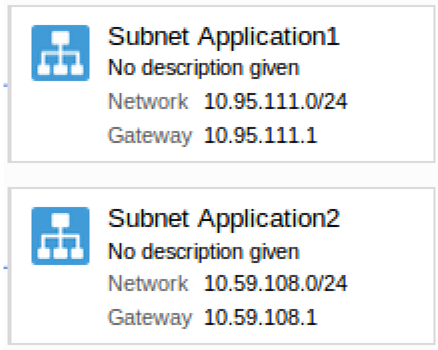

Subnet Application1
No description given
Network 10.95.111.0/24
Gateway 10.95.111.1

Subnet Application2
No description given
Network 10.59.108.0/24
Gateway 10.59.108.1

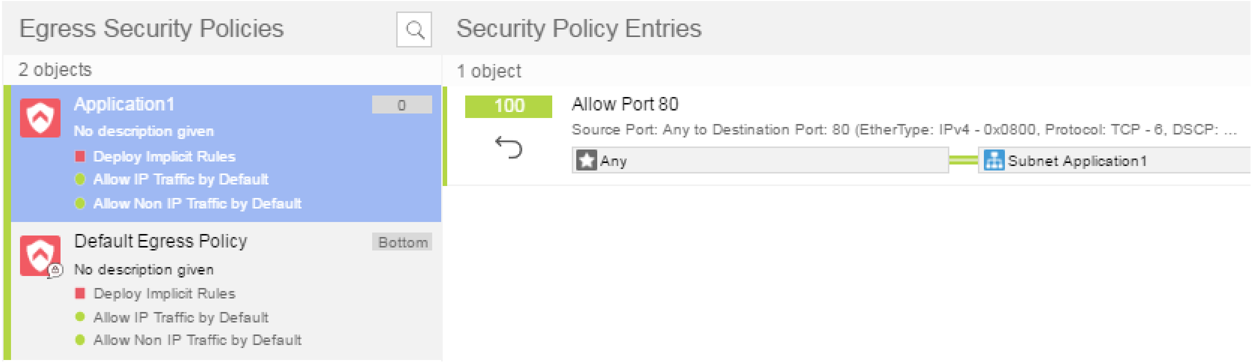

Egress Security Policies

2 objects

Application1 — 0

No description given

- 🟥 Deploy Implicit Rules
- 🟢 Allow IP Traffic by Default
- 🟢 Allow Non IP Traffic by Default

Default Egress Policy — Bottom

No description given

- 🟥 Deploy Implicit Rules
- 🟢 Allow IP Traffic by Default
- 🟢 Allow Non IP Traffic by Default

Security Policy Entries

1 object

100 — Allow Port 80

Source Port: Any to Destination Port: 80 (EtherType: IPv4 - 0x0800, Protocol: TCP - 6, DSCP: ...

Any — Subnet Application1

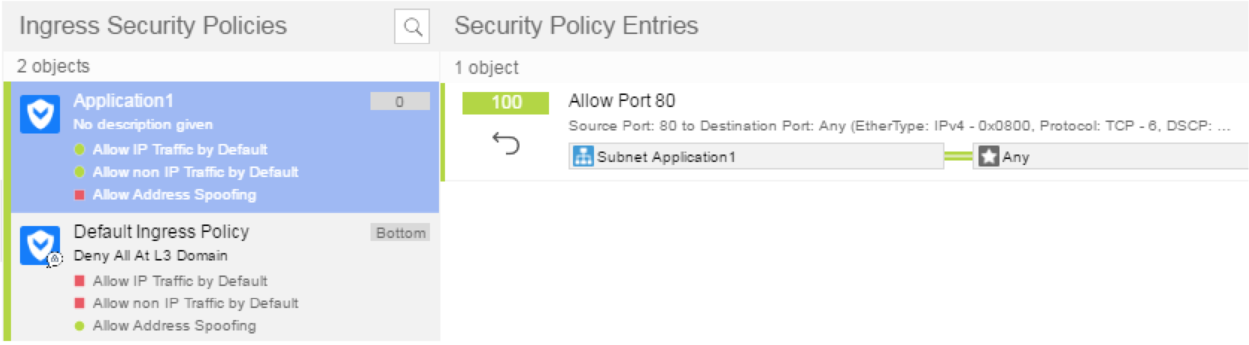

Ingress Security Policies

2 objects

Application1 — 0

No description given

- 🟢 Allow IP Traffic by Default
- 🟢 Allow non IP Traffic by Default
- 🟥 Allow Address Spoofing

Default Ingress Policy — Bottom

Deny All At L3 Domain

- 🟥 Allow IP Traffic by Default
- 🟥 Allow non IP Traffic by Default
- 🟢 Allow Address Spoofing

Security Policy Entries

1 object

100 — Allow Port 80

Source Port: 80 to Destination Port: Any (EtherType: IPv4 - 0x0800, Protocol: TCP - 6, DSCP: ...

Subnet Application1 — Any

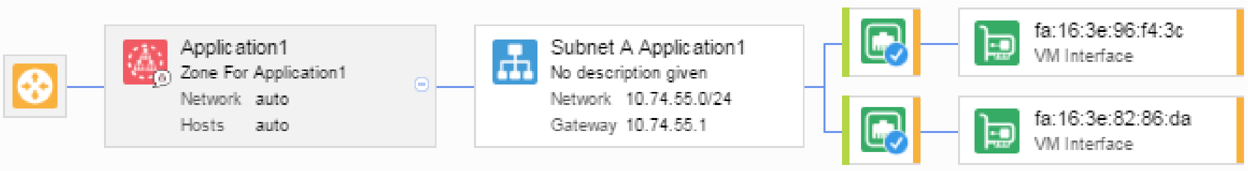

Application1
Zone For Application1
Network auto
Hosts auto

Subnet A Application1
No description given
Network 10.74.55.0/24
Gateway 10.74.55.1

fa:16:3e:96:f4:3c
VM Interface

fa:16:3e:82:86:da
VM Interface

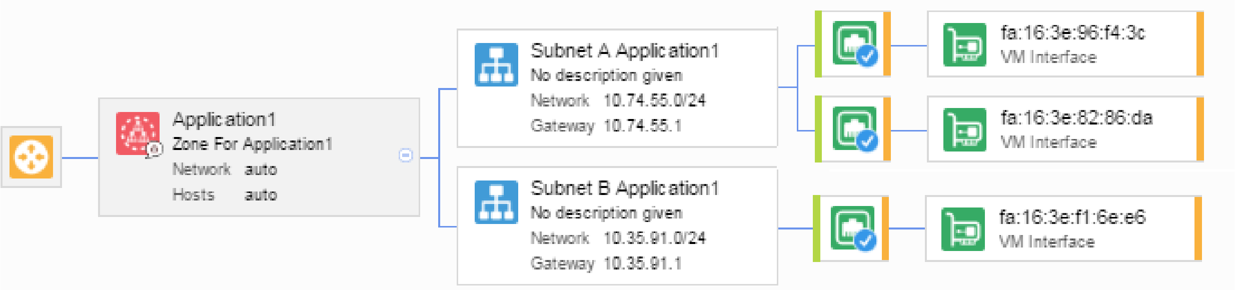

Application1
Zone For Application1
Network auto
Hosts auto

Subnet A Application1
No description given
Network 10.74.55.0/24
Gateway 10.74.55.1

Subnet B Application1
No description given
Network 10.35.91.0/24
Gateway 10.35.91.1

fa:16:3e:96:f4:3c
VM Interface

fa:16:3e:82:86:da
VM Interface

fa:16:3e:f1:6e:e6
VM Interface

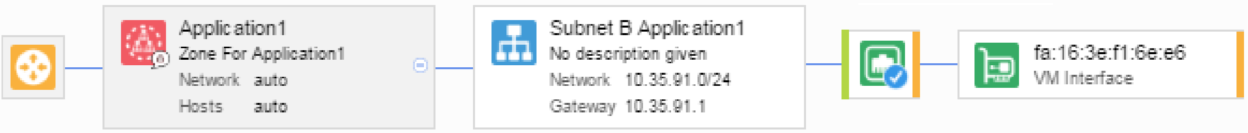

Application1
Zone For Application1
Network auto
Hosts auto

Subnet B Application1
No description given
Network 10.35.91.0/24
Gateway 10.35.91.1

fa:16:3e:f1:6e:e6
VM Interface

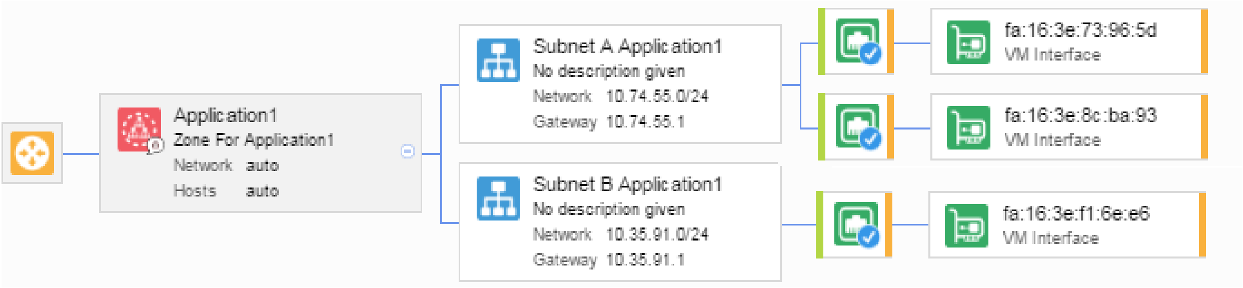

Application1
Zone For Application1
Network auto
Hosts auto

Subnet A Application1
No description given
Network 10.74.55.0/24
Gateway 10.74.55.1

Subnet B Application1
No description given
Network 10.35.91.0/24
Gateway 10.35.91.1

fa:16:3e:73:96:5d
VM Interface

fa:16:3e:8c:ba:93
VM Interface

fa:16:3e:f1:6e:e6
VM Interface

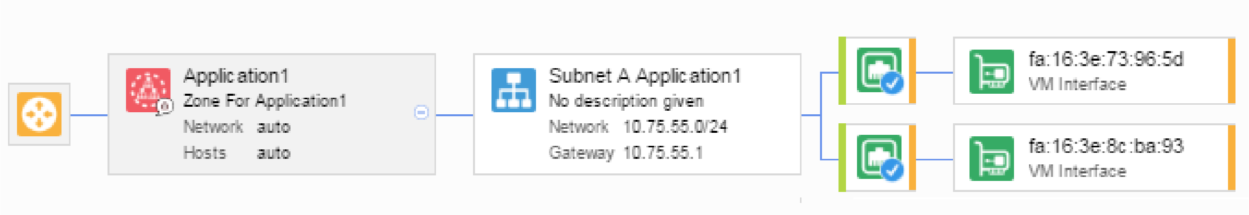

Application1
Zone For Application1
Network auto
Hosts auto

Subnet A Application1
No description given
Network 10.75.55.0/24
Gateway 10.75.55.1

fa:16:3e:73:96:5d
VM Interface

fa:16:3e:8c:ba:93
VM Interface

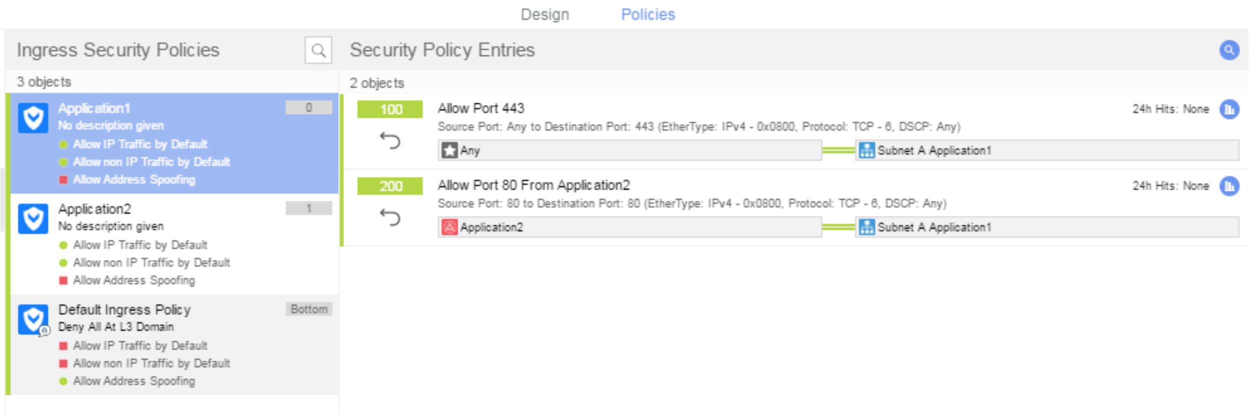

Ingress Security Policies

3 objects

Application1 0
No description given
● Allow IP Traffic by Default
● Allow non IP Traffic by Default
■ Allow Address Spoofing

Application2 1
No description given
● Allow IP Traffic by Default
● Allow non IP Traffic by Default
■ Allow Address Spoofing

Default Ingress Policy Bottom
Deny All At L3 Domain
■ Allow IP Traffic by Default
■ Allow non IP Traffic by Default
● Allow Address Spoofing

Security Policy Entries

2 objects

100 Allow Port 443 24h Hits: None
Source Port: Any to Destination Port: 443 (EtherType: IPv4 - 0x0800, Protocol: TCP - 6, DSCP: Any)
Any — Subnet A Application1

200 Allow Port 80 From Application2 24h Hits: None
Source Port: 80 to Destination Port: 80 (EtherType: IPv4 - 0x0800, Protocol: TCP - 6, DSCP: Any)
Application2 — Subnet A Application1

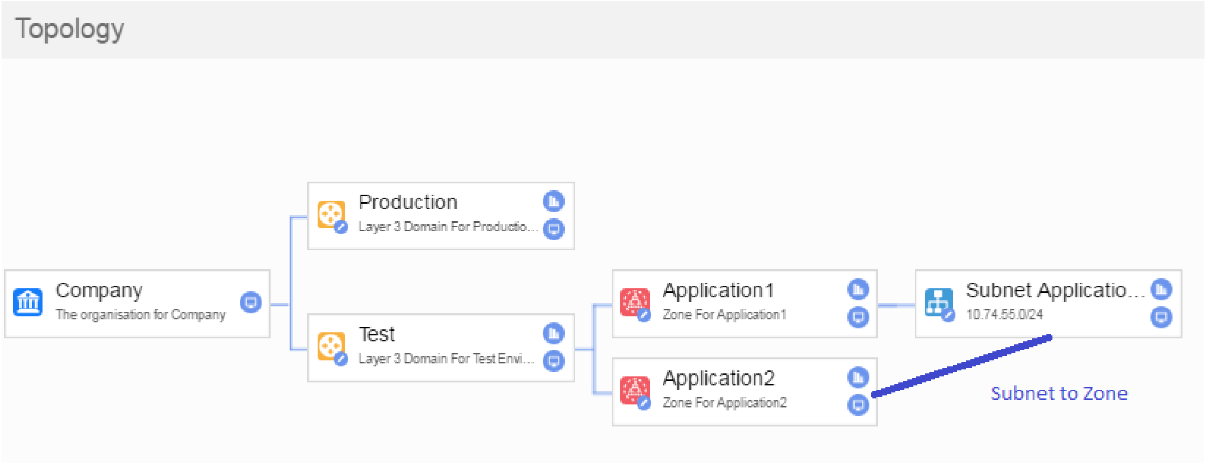

Topology

Company
The organisation for Company

Production
Layer 3 Domain For Productio...

Test
Layer 3 Domain For Test Envi...

Application1
Zone For Application1

Application2
Zone For Application2

Subnet Applicatio...
10.74.55.0/24

Subnet to Zone

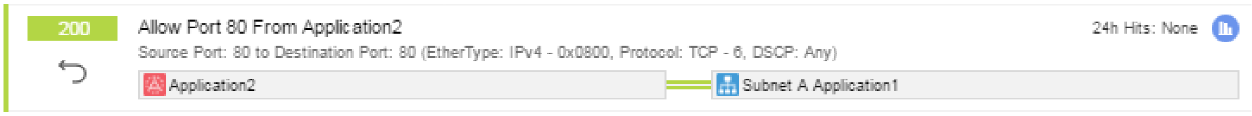

200 Allow Port 80 From Application2 24h Hits: None
Source Port: 80 to Destination Port: 80 (EtherType: IPv4 - 0x0800, Protocol: TCP - 6, DSCP: Any)
Application2 — Subnet A Application1

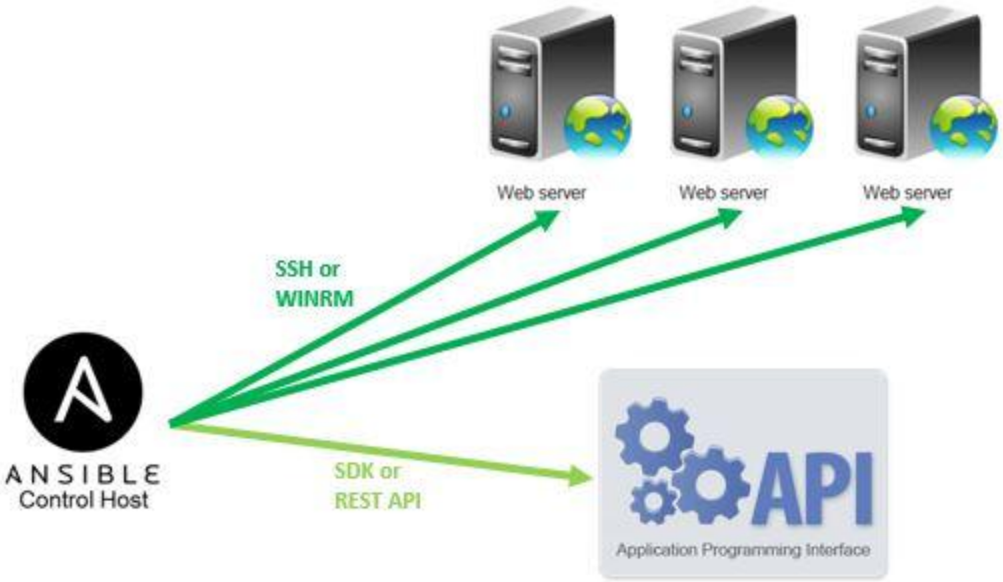
Web server
Web server
Web server
SSH or
WINRM
ANSIBLE
Control Host
SDK or
REST API
API
Application Programming Interface

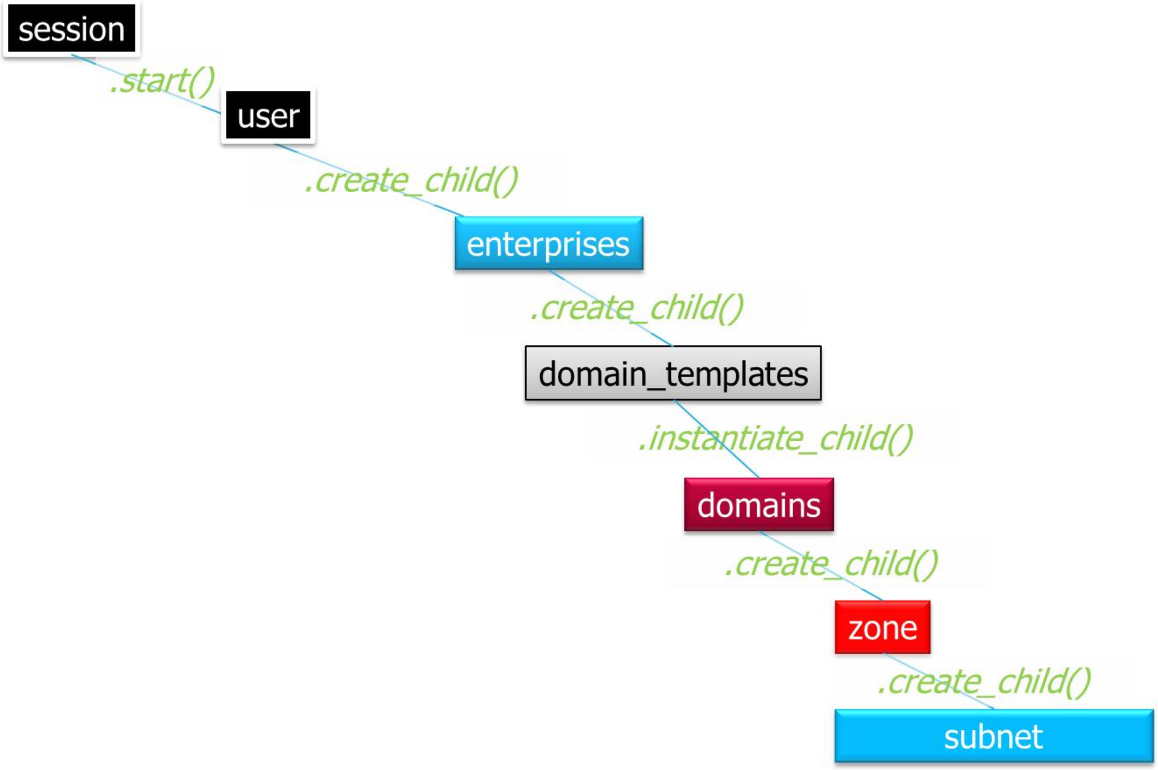
session
.start()
user
.create_child()
enterprises
.create_child()
domain_templates
.instantiate_child()
domains
.create_child()
zone
.create_child()
subnet

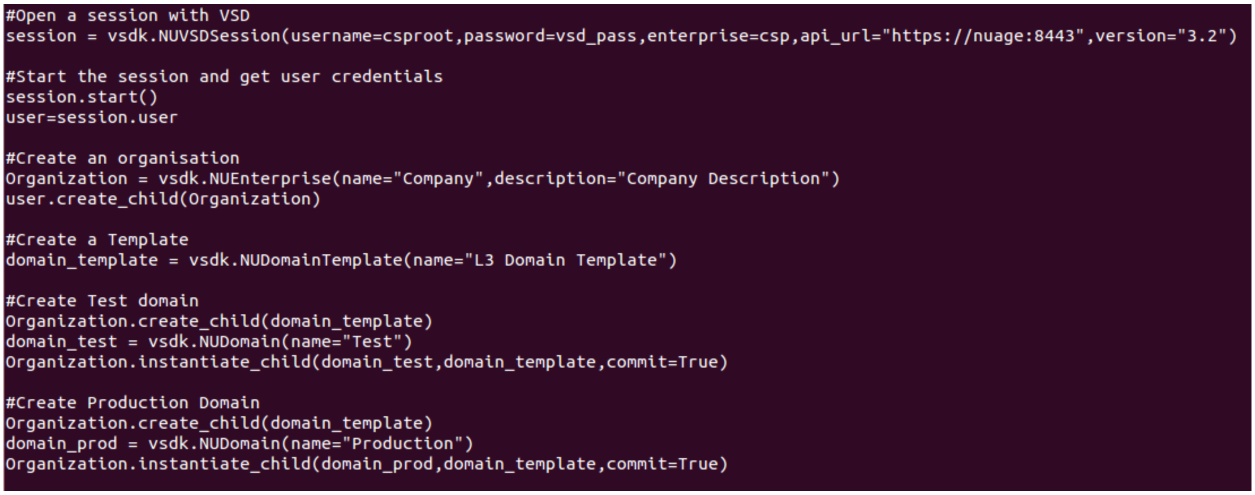

```
#Open a session with VSD
session = vsdk.NUVSDSession(username=csproot,password=vsd_pass,enterprise=csp,api_url="https://nuage:8443",version="3.2")

#Start the session and get user credentials
session.start()
user=session.user

#Create an organisation
Organization = vsdk.NUEnterprise(name="Company",description="Company Description")
user.create_child(Organization)

#Create a Template
domain_template = vsdk.NUDomainTemplate(name="L3 Domain Template")

#Create Test domain
Organization.create_child(domain_template)
domain_test = vsdk.NUDomain(name="Test")
Organization.instantiate_child(domain_test,domain_template,commit=True)

#Create Production Domain
Organization.create_child(domain_template)
domain_prod = vsdk.NUDomain(name="Production")
Organization.instantiate_child(domain_prod,domain_template,commit=True)
```

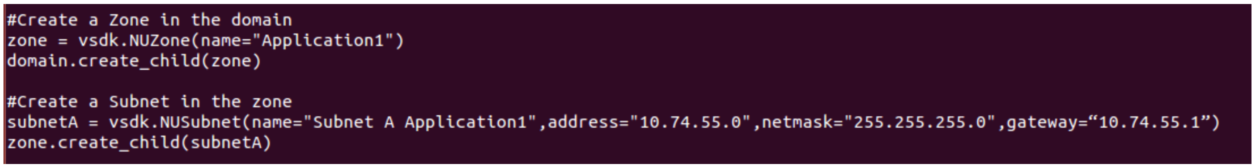

```
#Create a Zone in the domain
zone = vsdk.NUZone(name="Application1")
domain.create_child(zone)

#Create a Subnet in the zone
subnetA = vsdk.NUSubnet(name="Subnet A Application1",address="10.74.55.0",netmask="255.255.255.0",gateway=“10.74.55.1”)
zone.create_child(subnetA)
```

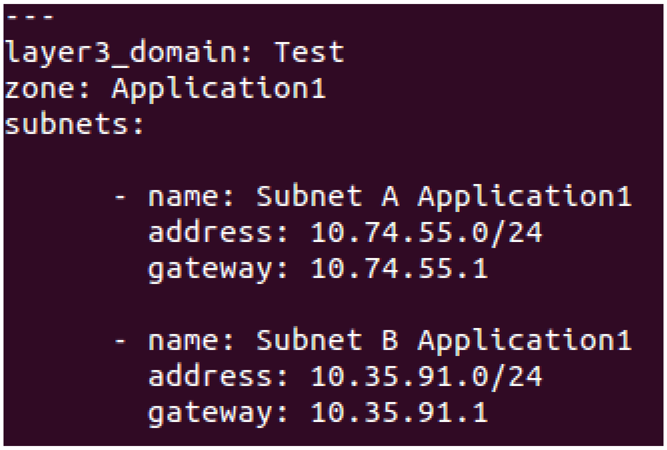

```
---
layer3_domain: Test
zone: Application1
subnets:

      - name: Subnet A Application1
        address: 10.74.55.0/24
        gateway: 10.74.55.1

      - name: Subnet B Application1
        address: 10.35.91.0/24
        gateway: 10.35.91.1
```

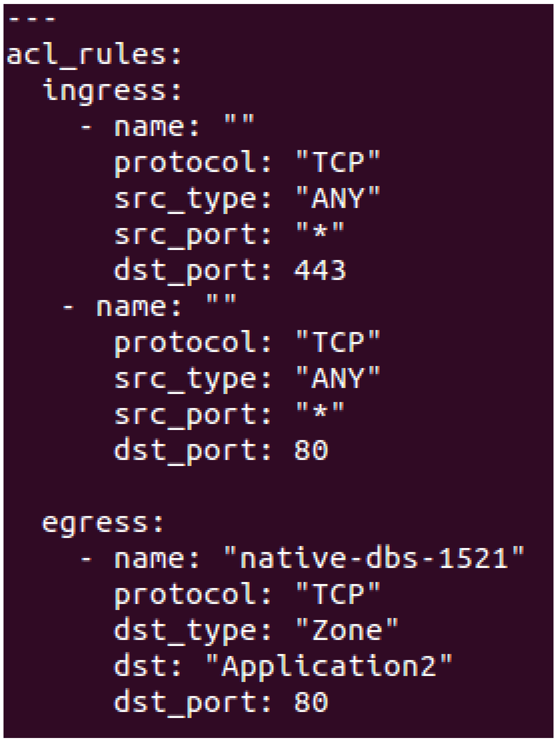

```
---
acl_rules:
  ingress:
    - name: ""
      protocol: "TCP"
      src_type: "ANY"
      src_port: "*"
      dst_port: 443
   - name: ""
      protocol: "TCP"
      src_type: "ANY"
      src_port: "*"
      dst_port: 80

  egress:
    - name: "native-dbs-1521"
      protocol: "TCP"
      dst_type: "Zone"
      dst: "Application2"
      dst_port: 80
```

Chapter 7: Using Continuous Integration Builds For Network Configuration

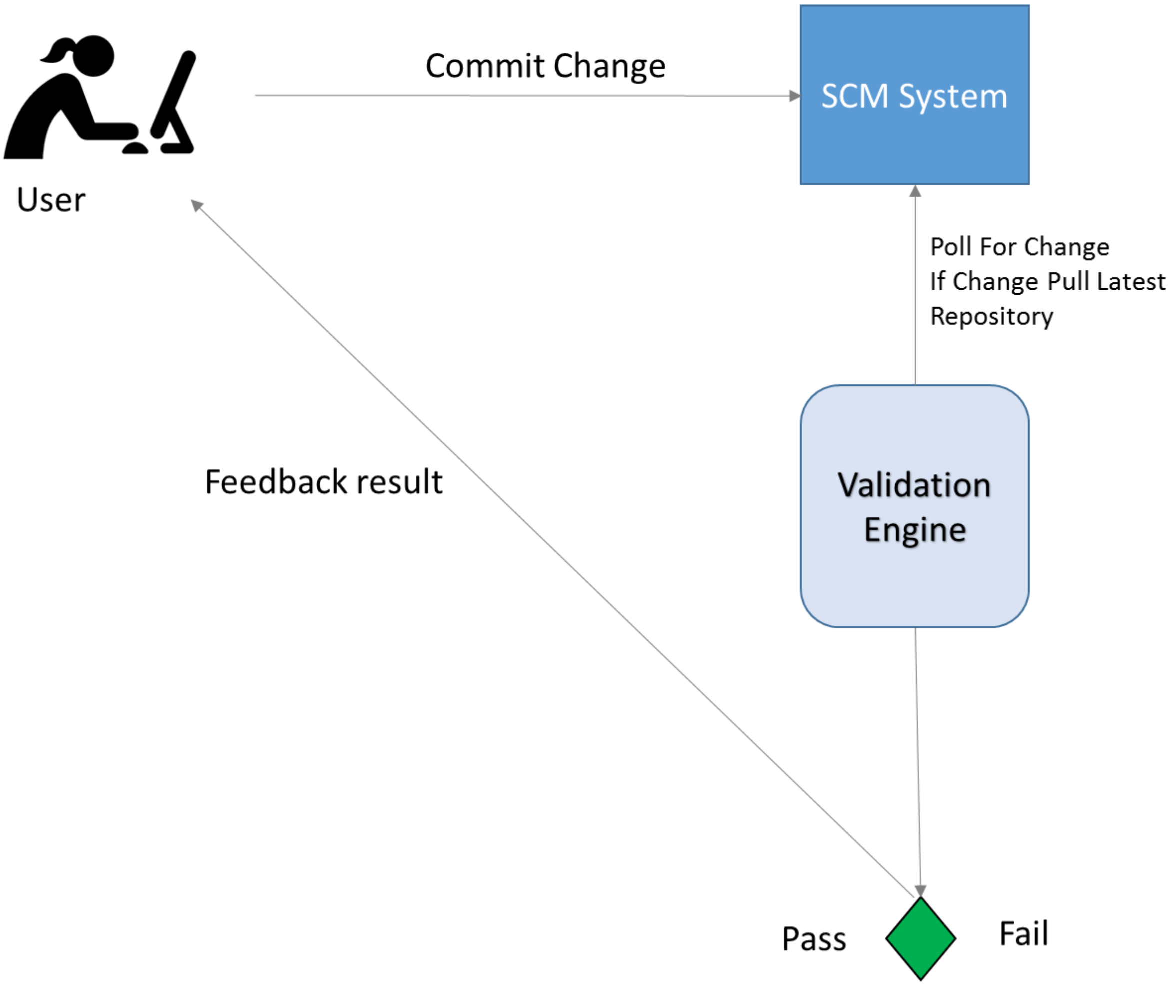

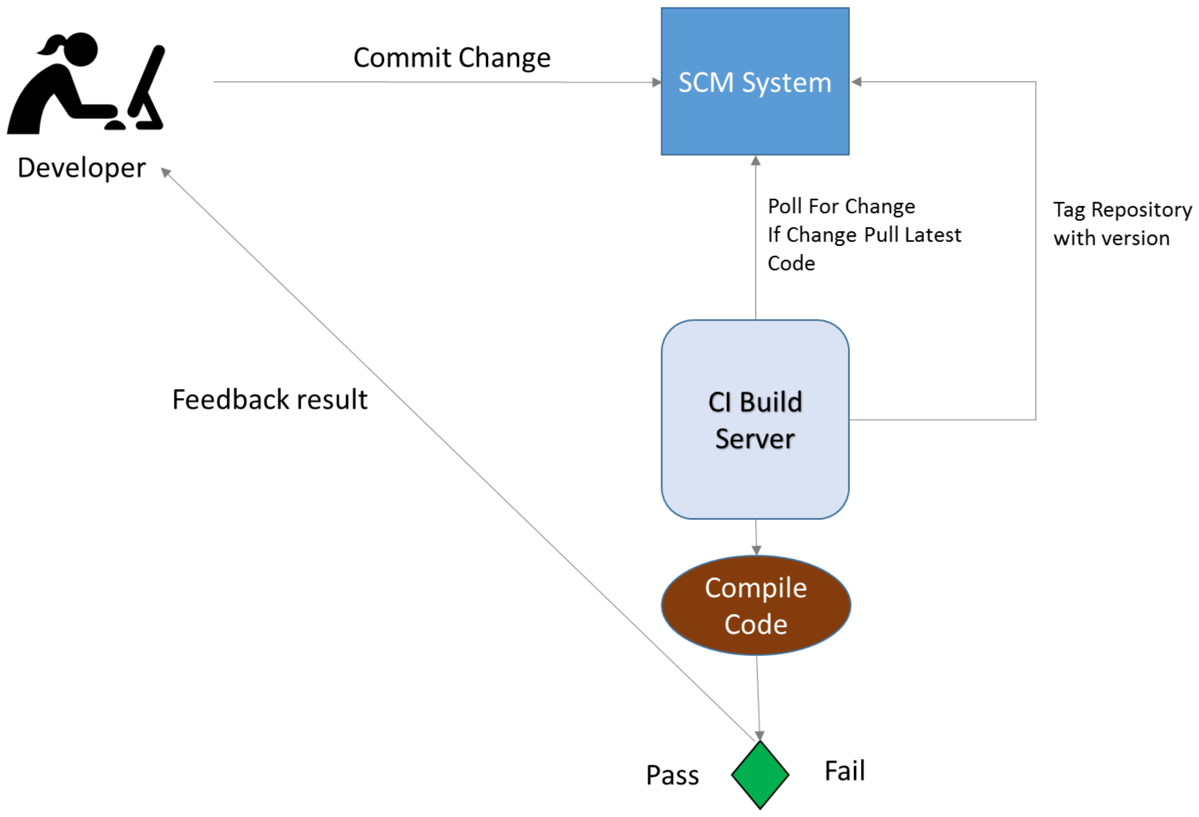
Commit Change
SCM System
Developer
Poll For Change
If Change Pull Latest
Code
Tag Repository
with version
Feedback result
CI Build
Server
Compile
Code
Pass
Fail

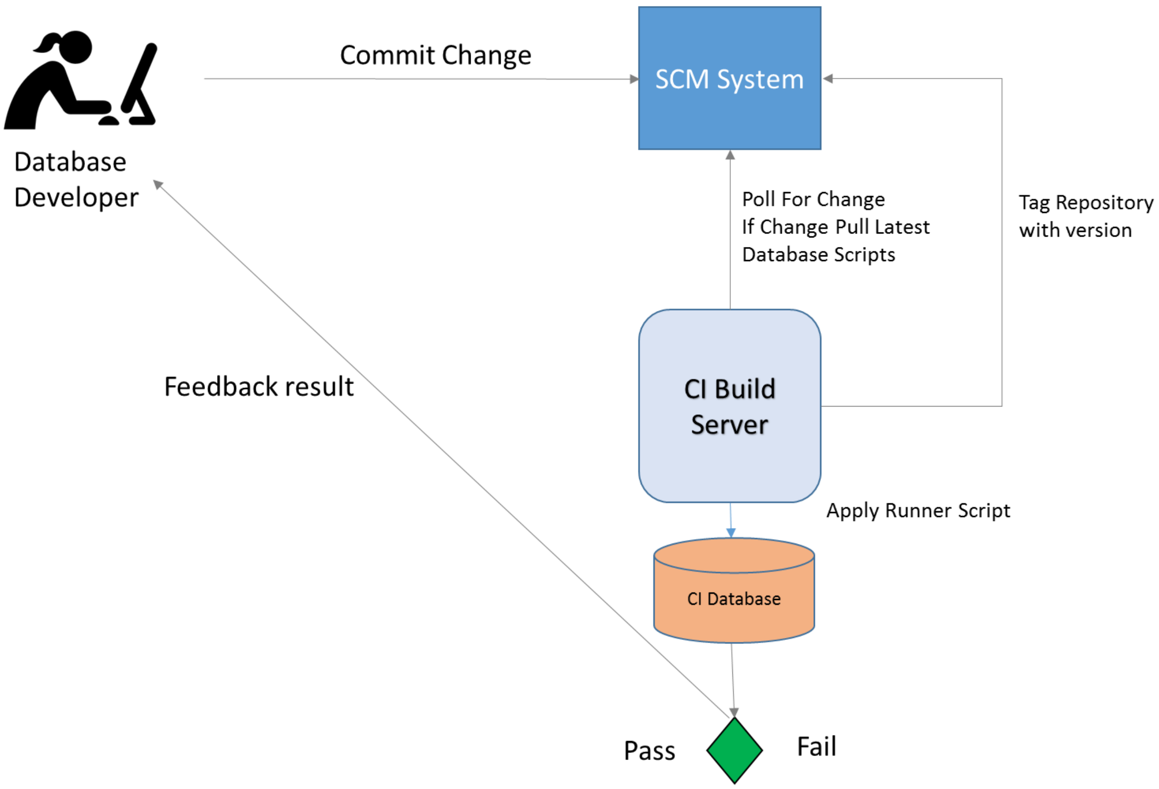
Commit Change
SCM System
Database
Developer
Poll For Change
If Change Pull Latest
Database Scripts
Tag Repository
with version
Feedback result
CI Build
Server
Apply Runner Script
CI Database
Pass
Fail

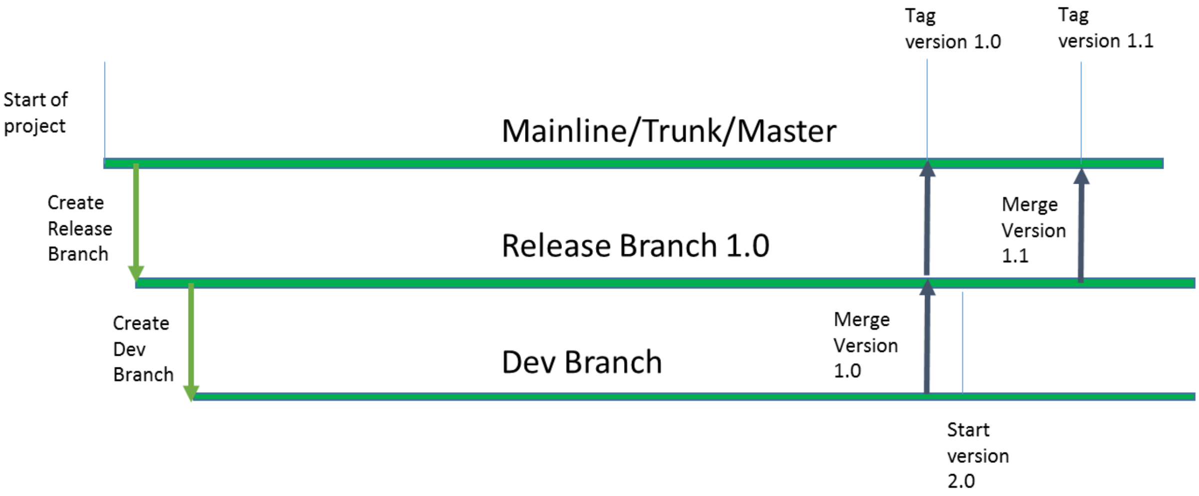
Tag
version 1.0
Tag
version 1.1
Start of
project
Mainline/Trunk/Master
Create
Release
Branch
Release Branch 1.0
Merge
Version
1.1
Create
Dev
Branch
Merge
Version
1.0
Dev Branch
Start
version
2.0

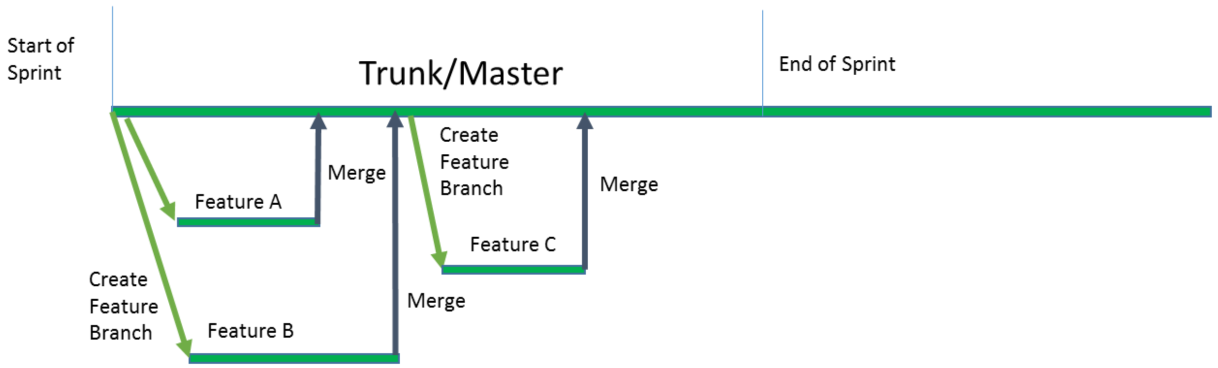
Start of Sprint
Trunk/Master
End of Sprint
Merge
Create Feature Branch
Merge
Feature A
Feature C
Create Feature Branch
Merge
Feature B

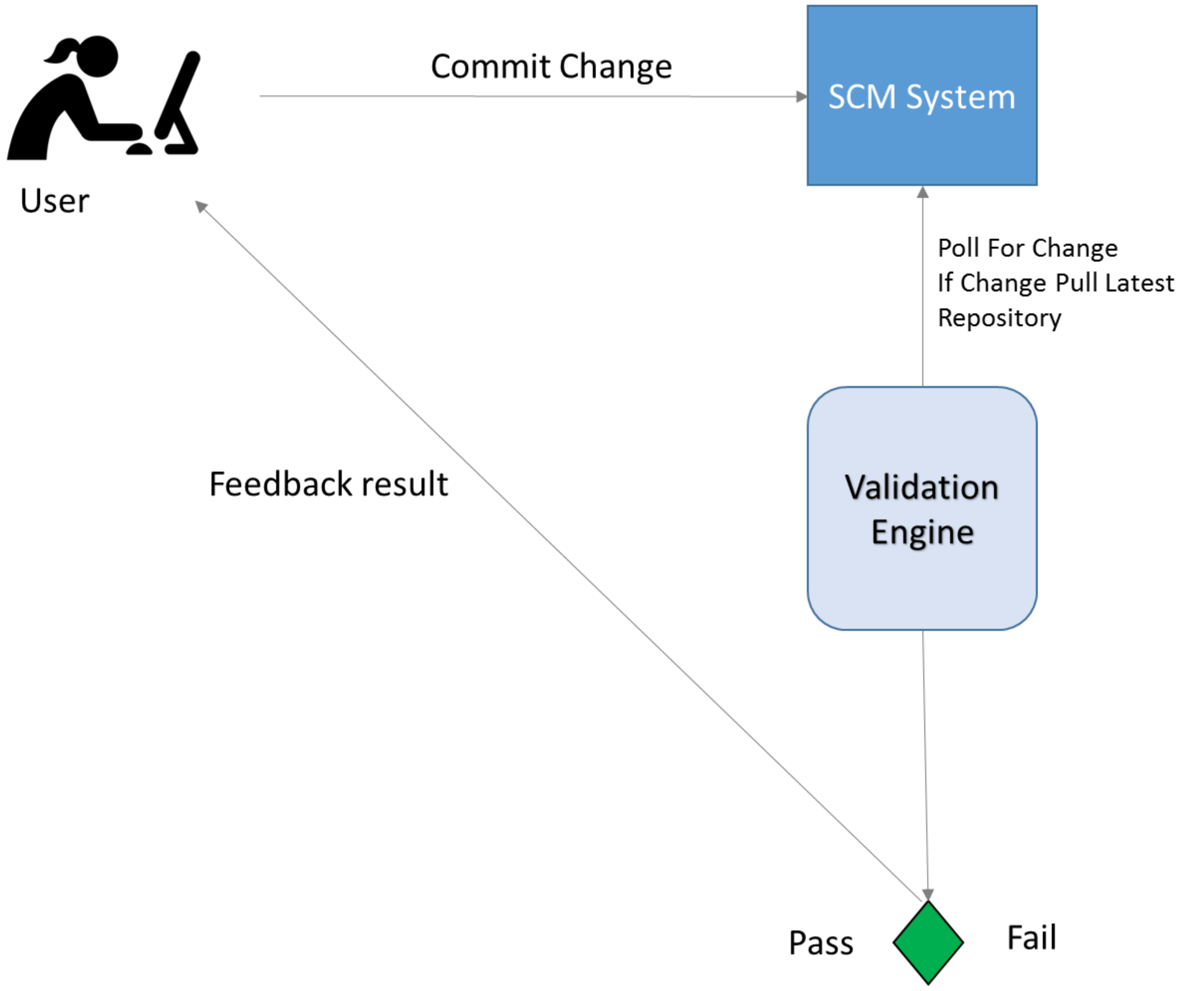
Commit Change
SCM System
User
Poll For Change
If Change Pull Latest Repository
Validation Engine
Feedback result
Pass
Fail

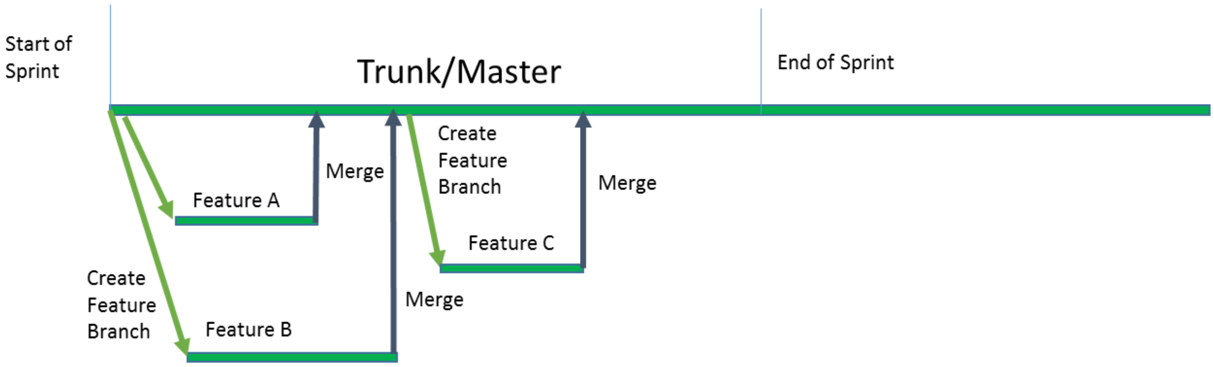

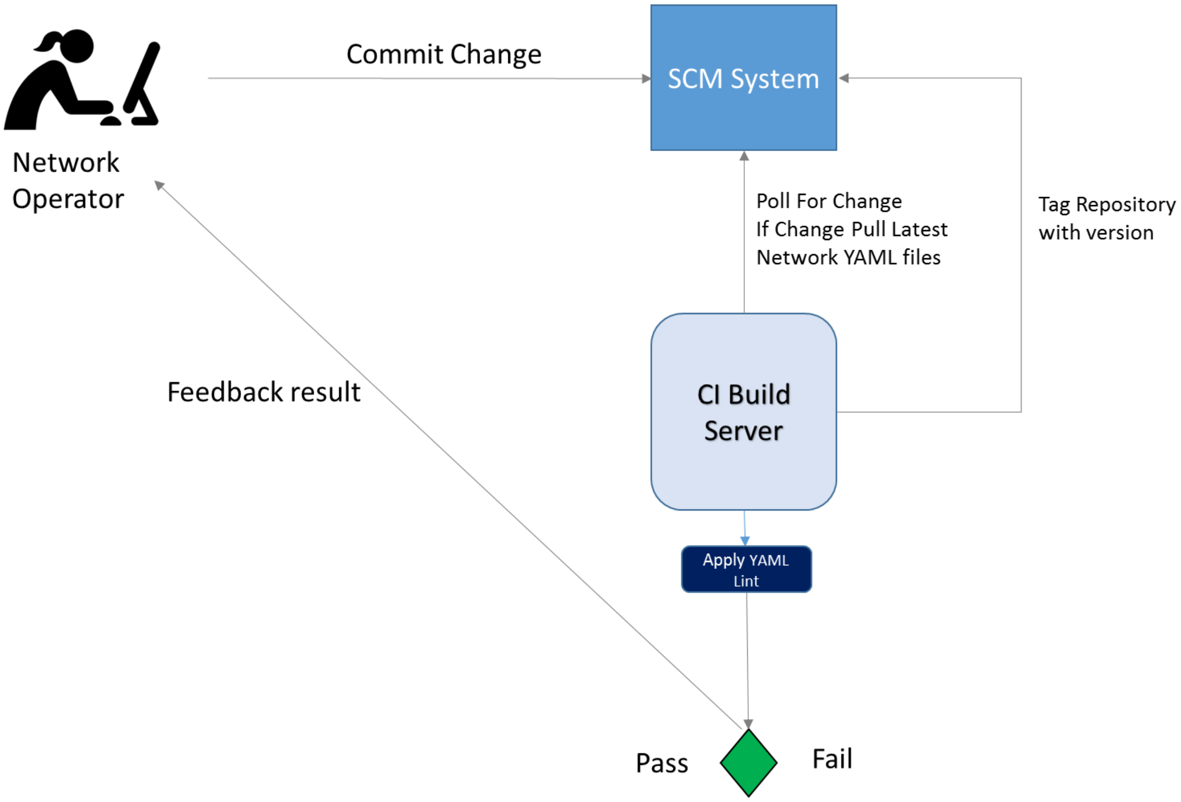

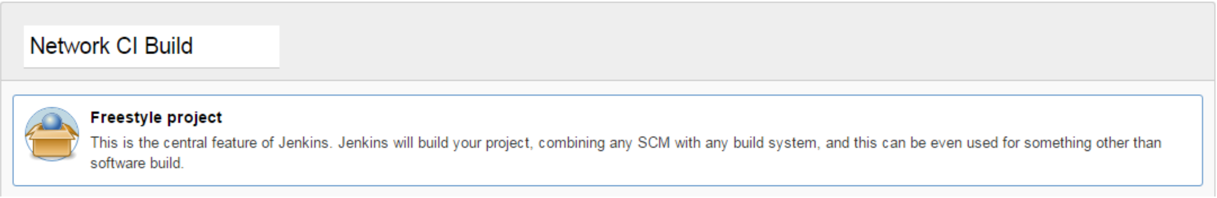
Network CI Build

Freestyle project

This is the central feature of Jenkins. Jenkins will build your project, combining any SCM with any build system, and this can be even used for something other than software build.

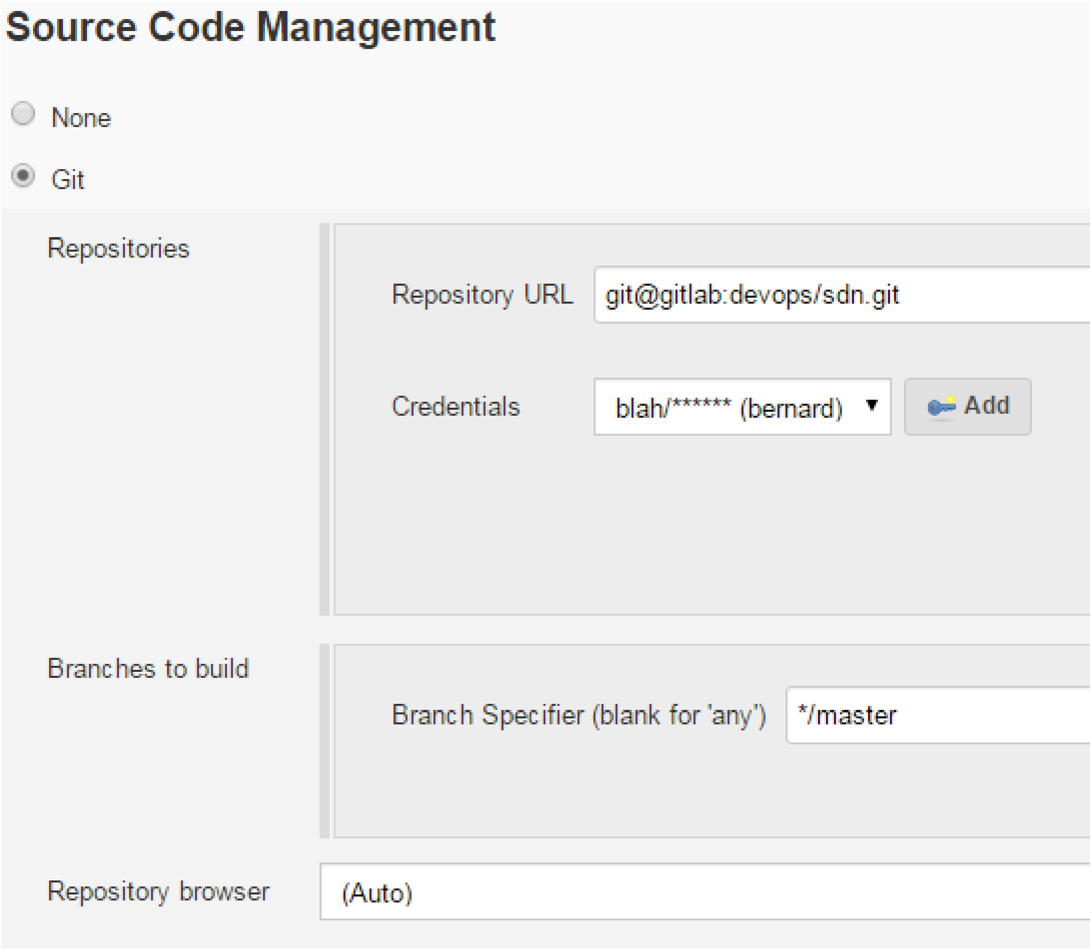

Source Code Management

○ None

● Git

Repositories

Repository URL: git@gitlab:devops/sdn.git

Credentials: blah/****** (bernard) ▾

Branches to build

Branch Specifier (blank for 'any'): */master

Repository browser: (Auto)

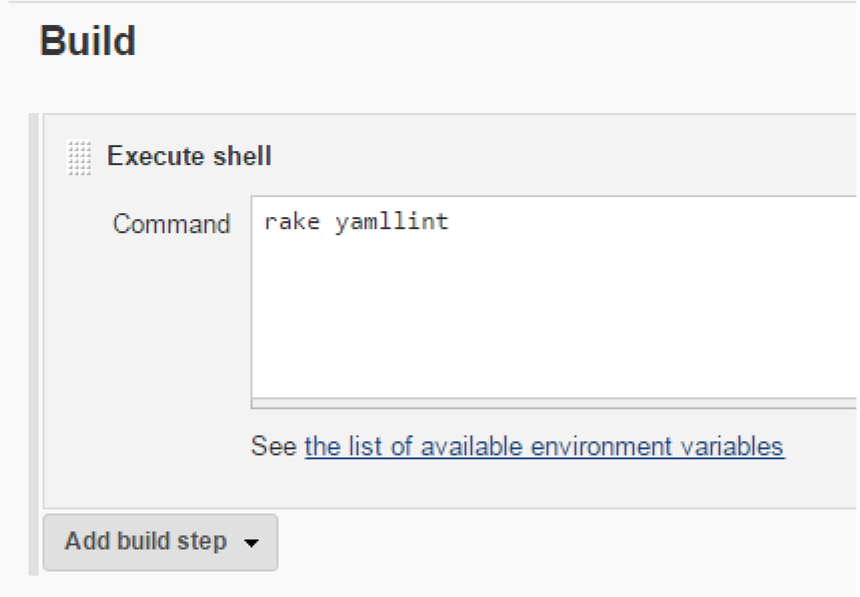

Build

Execute shell

Command:

```
rake yamllint
```

See the list of available environment variables

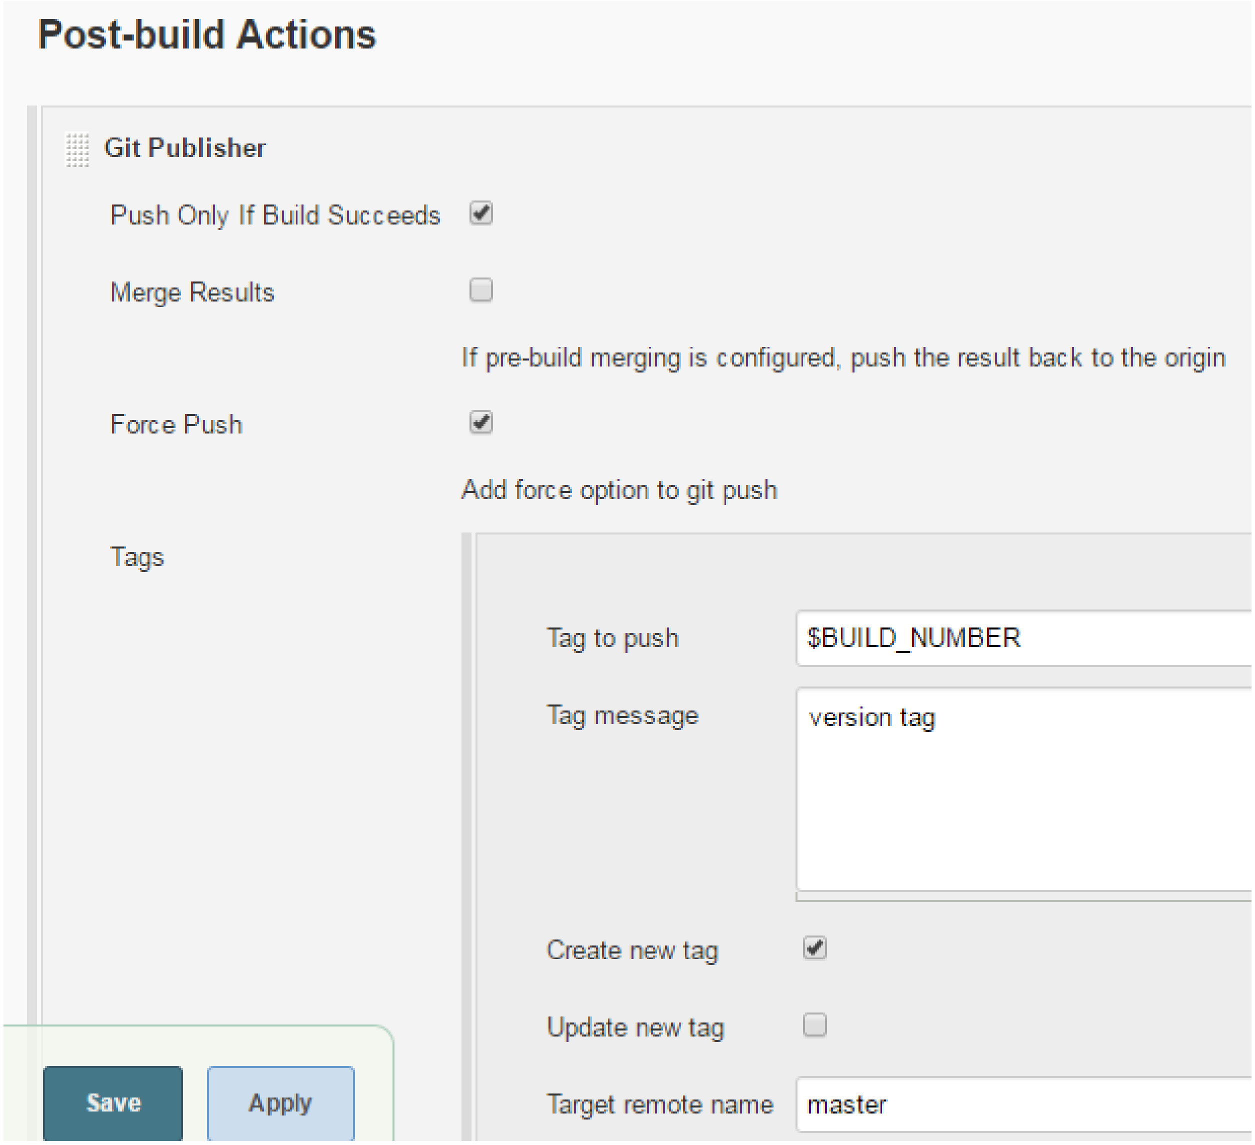

Post-build Actions

Git Publisher

Push Only If Build Succeeds ☑

Merge Results □

If pre-build merging is configured, push the result back to the origin

Force Push ☑

Add force option to git push

Tags

Tag to push: $BUILD_NUMBER

Tag message: version tag

Create new tag ☑

Update new tag □

Target remote name: master

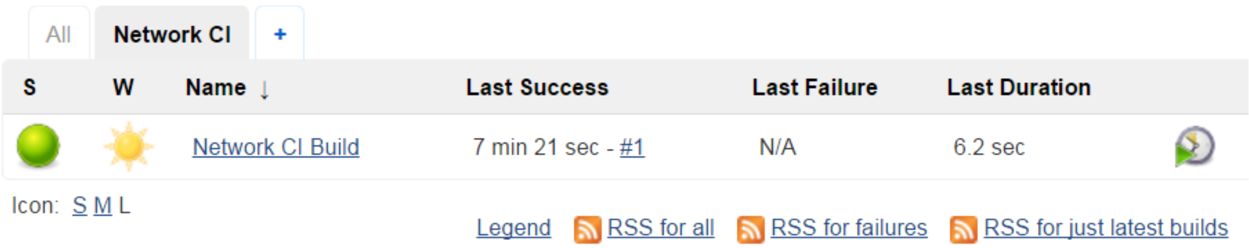

S	W	Name ↓	Last Success	Last Failure	Last Duration
●	☀	Network CI Build	7 min 21 sec - #1	N/A	6.2 sec

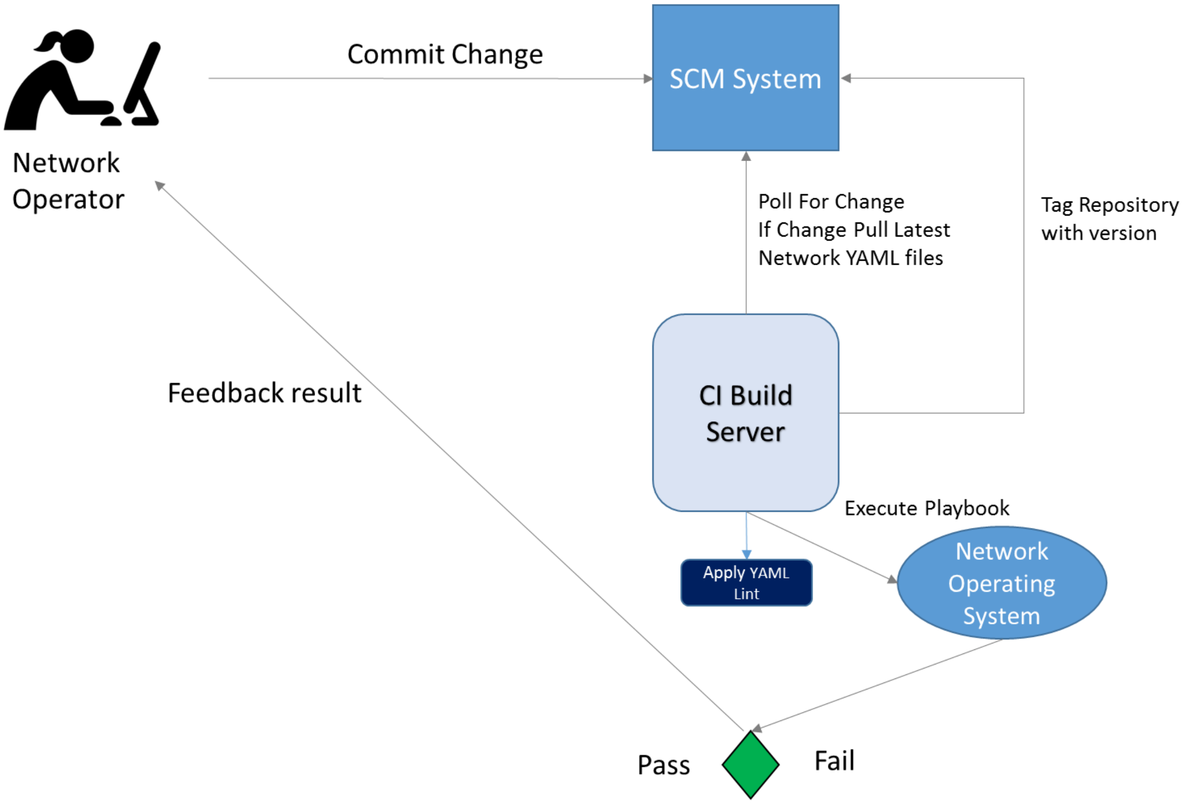
Network Operator
Commit Change
SCM System
Poll For Change
If Change Pull Latest
Network YAML files
Tag Repository
with version
Feedback result
CI Build Server
Execute Playbook
Apply YAML Lint
Network Operating System
Pass
Fail

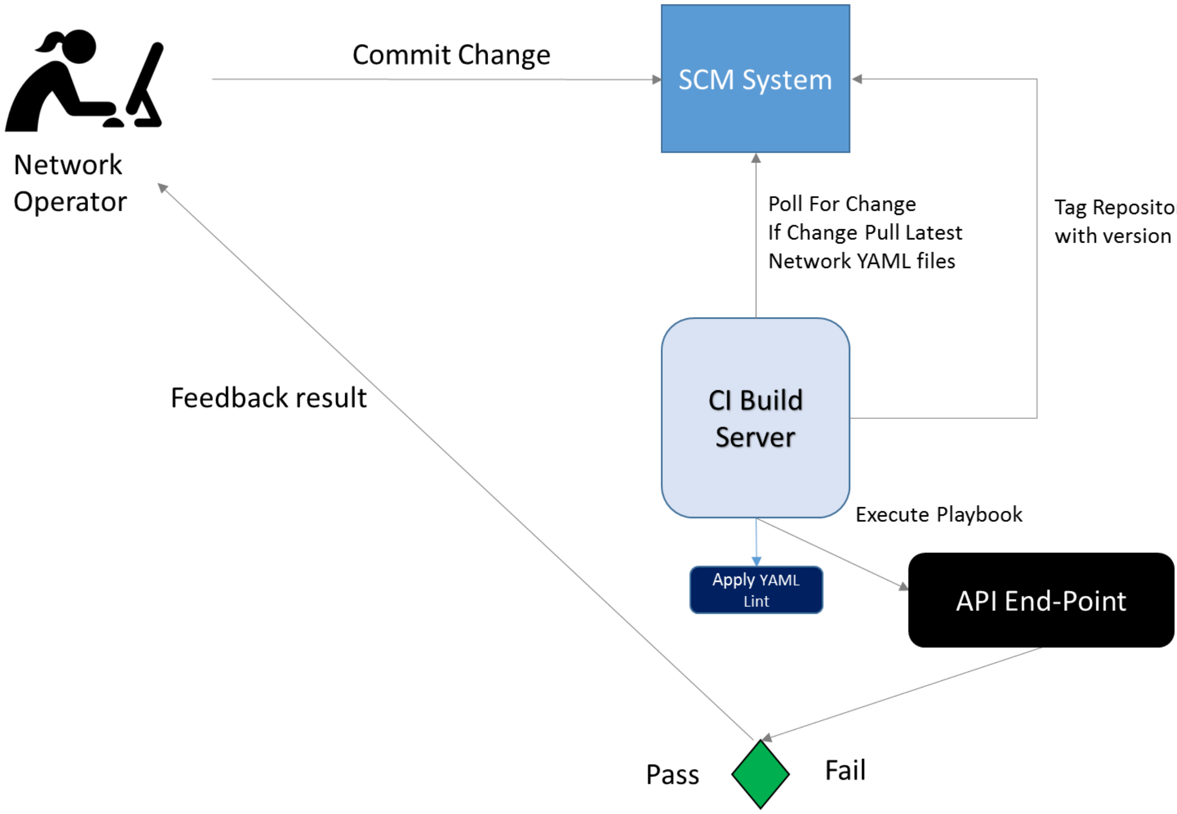
Commit Change
SCM System
Network Operator
Poll For Change
If Change Pull Latest
Network YAML files
Tag Repository
with version
Feedback result
CI Build Server
Execute Playbook
Apply YAML Lint
API End-Point
Pass
Fail

Chapter 8: Testing Network Changes

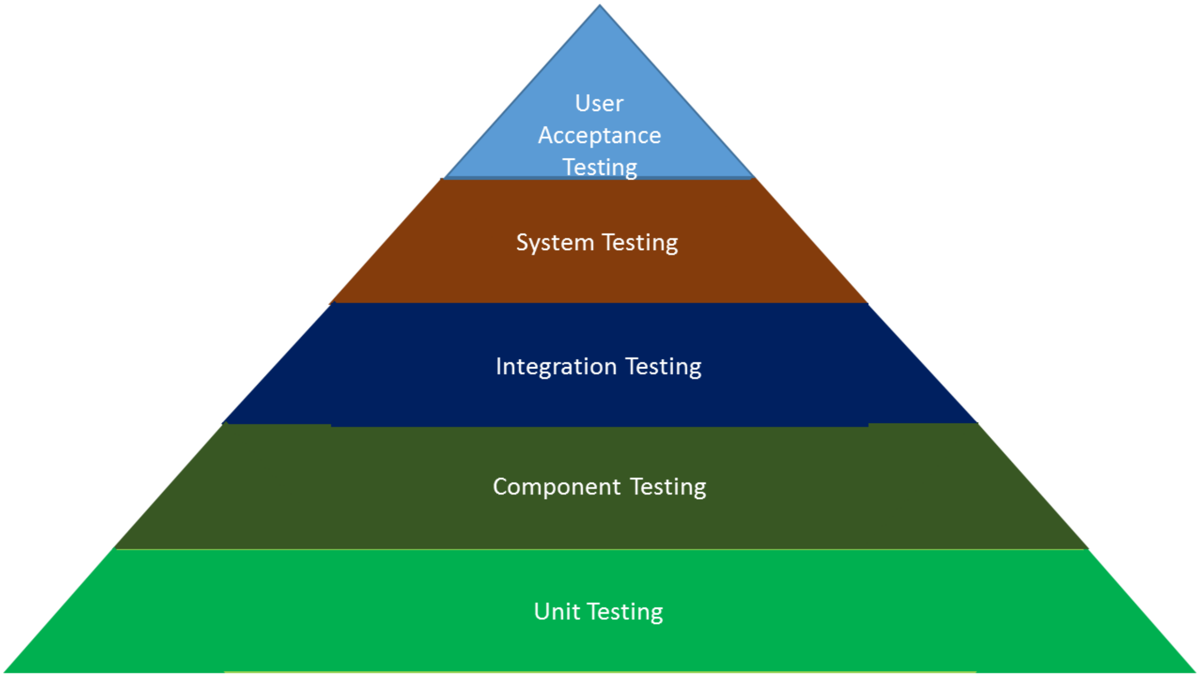

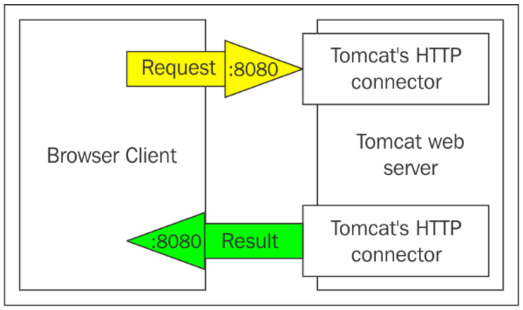

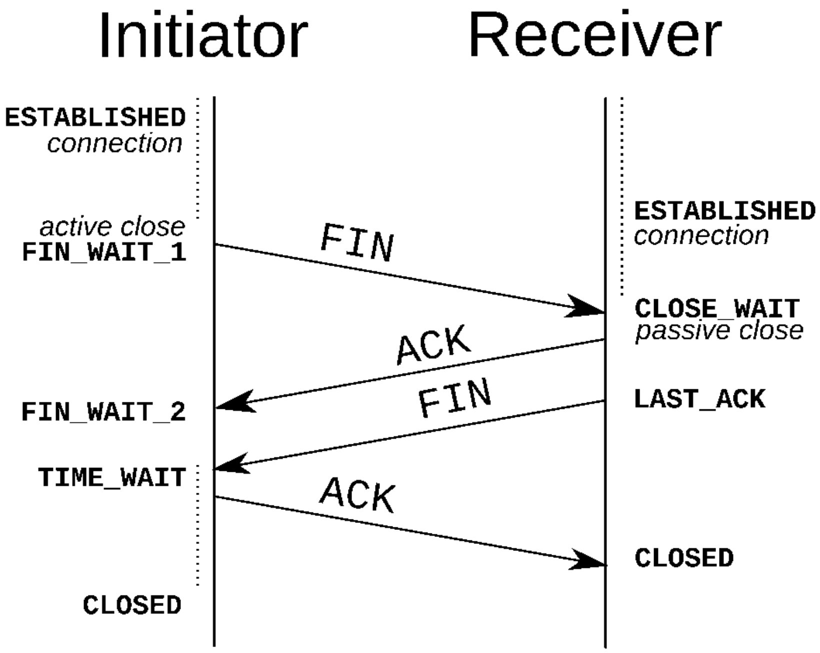
Initiator
Receiver
ESTABLISHED
connection
active close
FIN_WAIT_1
FIN
ESTABLISHED
connection
CLOSE_WAIT
passive close
ACK
FIN_WAIT_2
FIN
LAST_ACK
TIME_WAIT
ACK
CLOSED
CLOSED

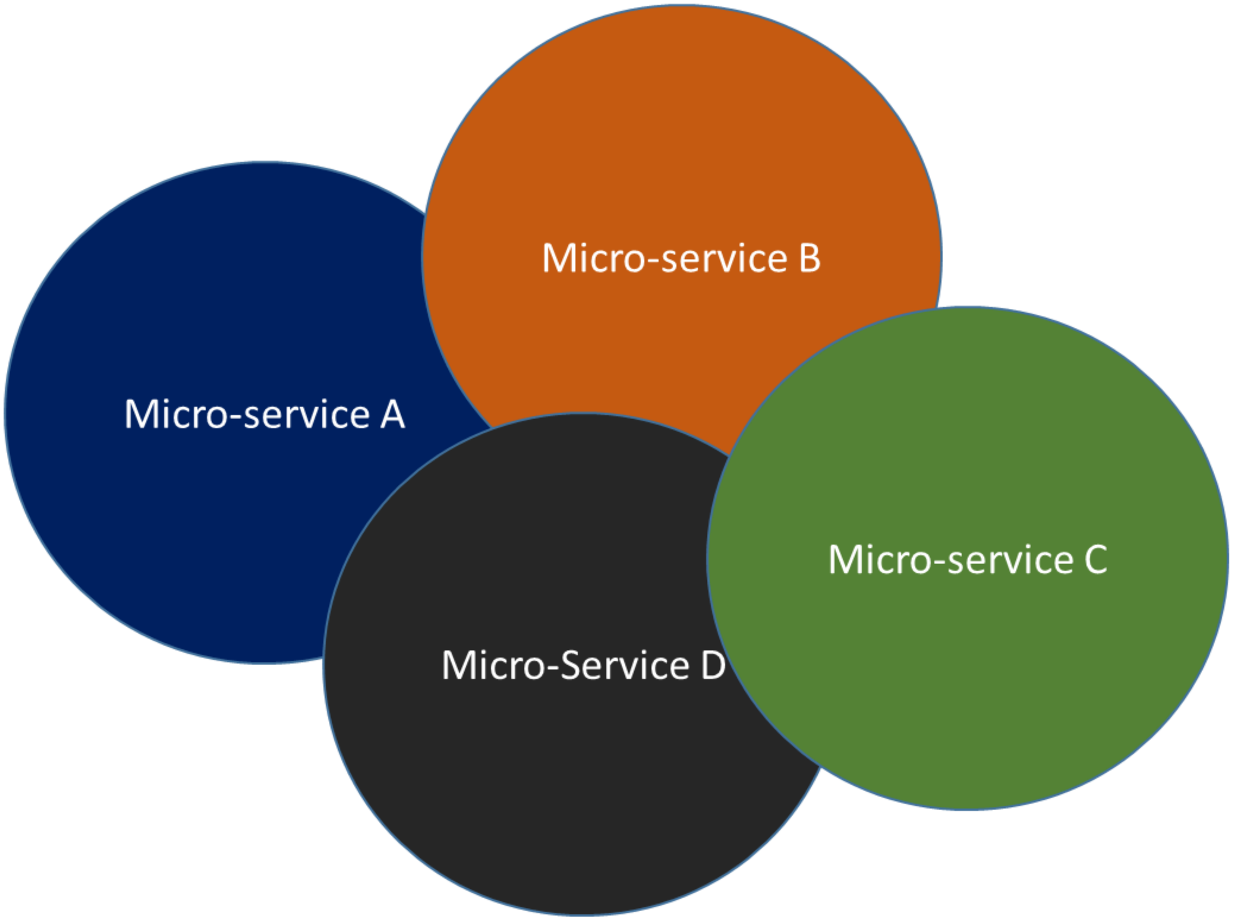
Micro-service B
Micro-service A
Micro-service C
Micro-Service D

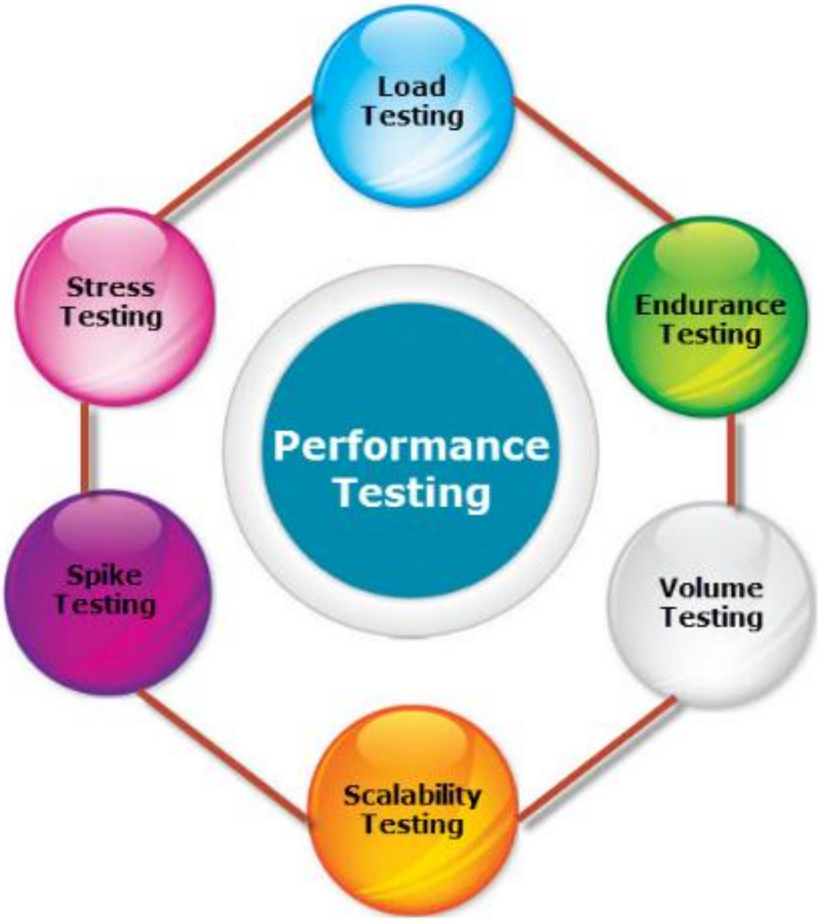
Load Testing
Stress Testing
Endurance Testing
Performance Testing
Spike Testing
Volume Testing
Scalability Testing

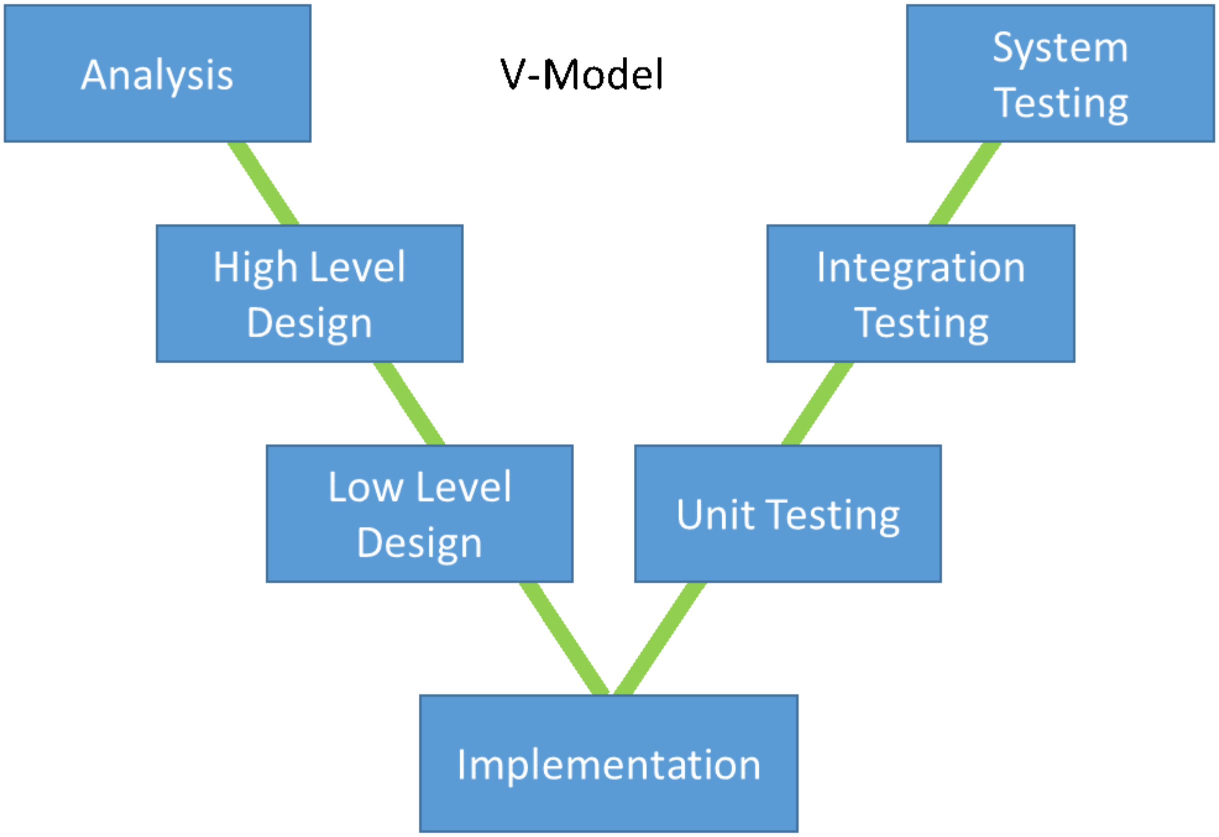
V-Model
Analysis
High Level Design
Low Level Design
Implementation
Unit Testing
Integration Testing
System Testing

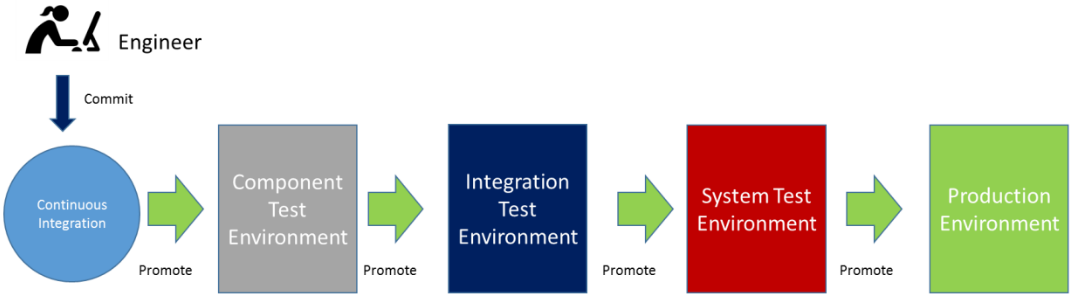
Engineer
Commit
Continuous Integration
Promote
Component Test Environment
Promote
Integration Test Environment
Promote
System Test Environment
Promote
Production Environment

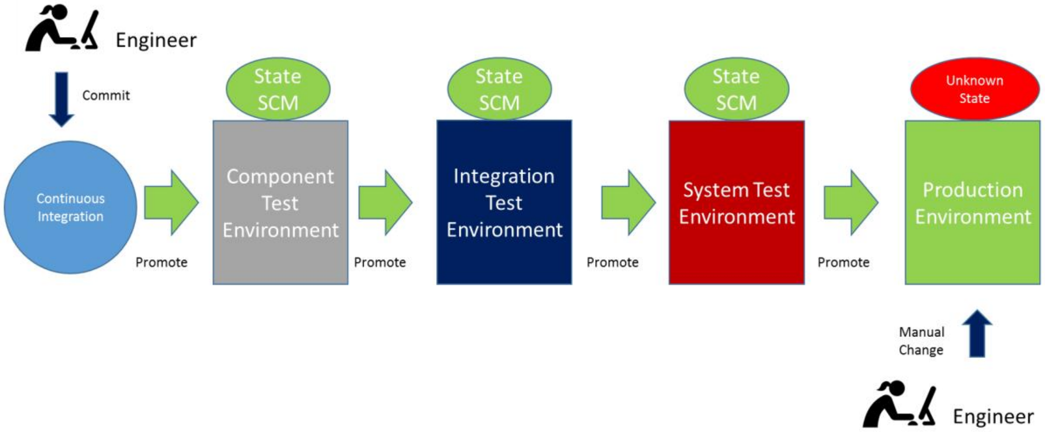
Engineer
Commit
Continuous Integration
Promote
State SCM
Component Test Environment
Promote
State SCM
Integration Test Environment
Promote
State SCM
System Test Environment
Promote
Unknown State
Production Environment
Manual Change
Engineer

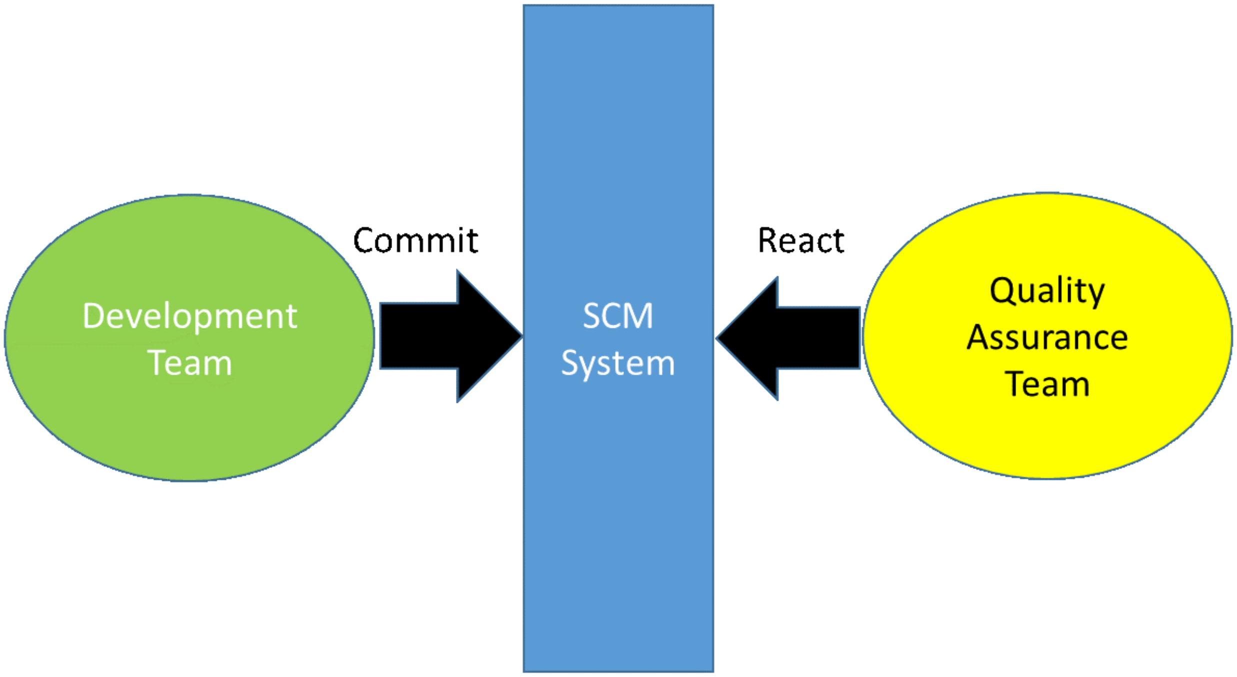
Development Team
Commit
SCM System
React
Quality Assurance Team

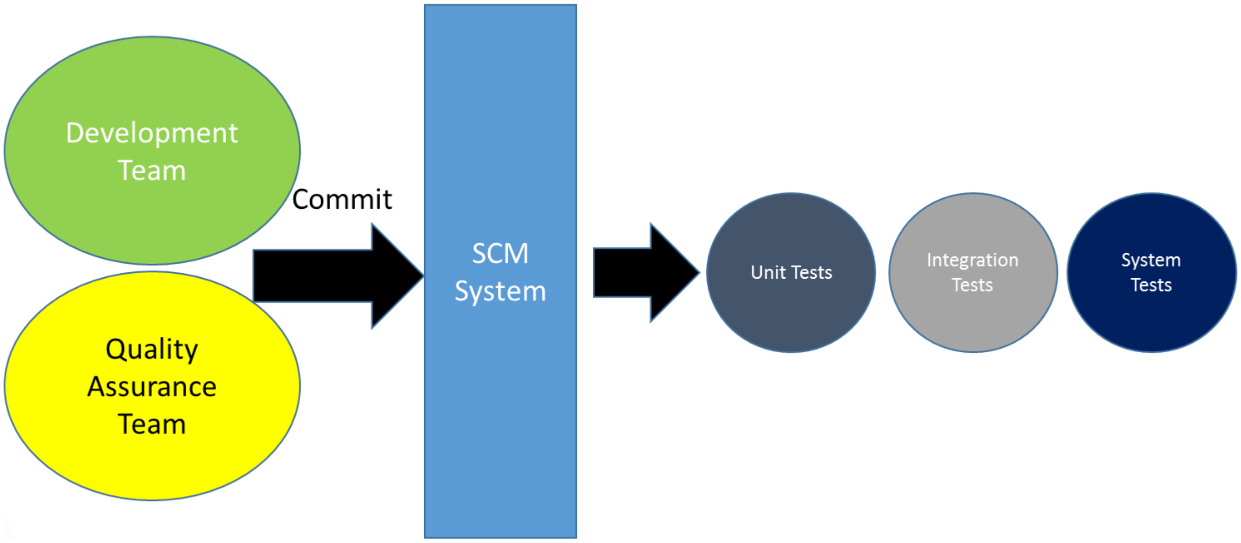
Development Team
Quality Assurance Team
Commit
SCM System
Unit Tests
Integration Tests
System Tests

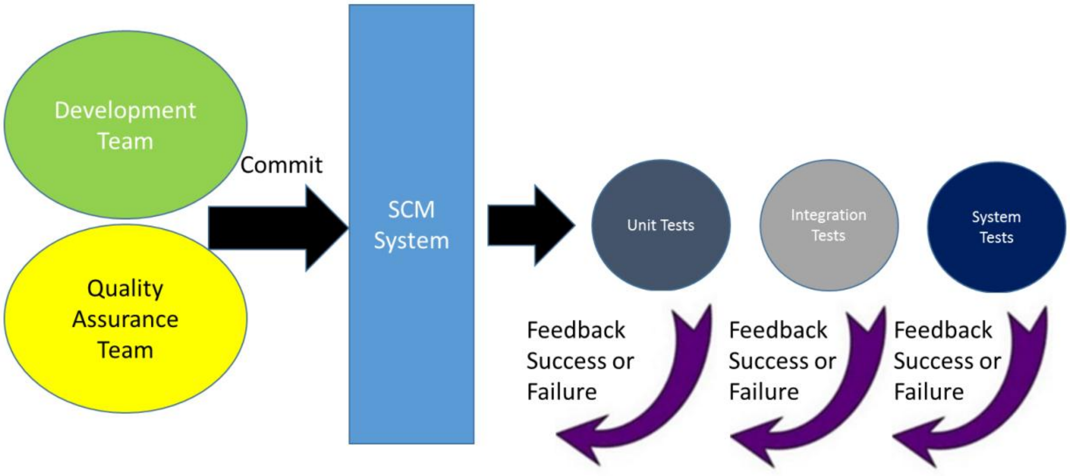
Development Team
Quality Assurance Team
Commit
SCM System
Unit Tests
Integration Tests
System Tests
Feedback Success or Failure
Feedback Success or Failure
Feedback Success or Failure

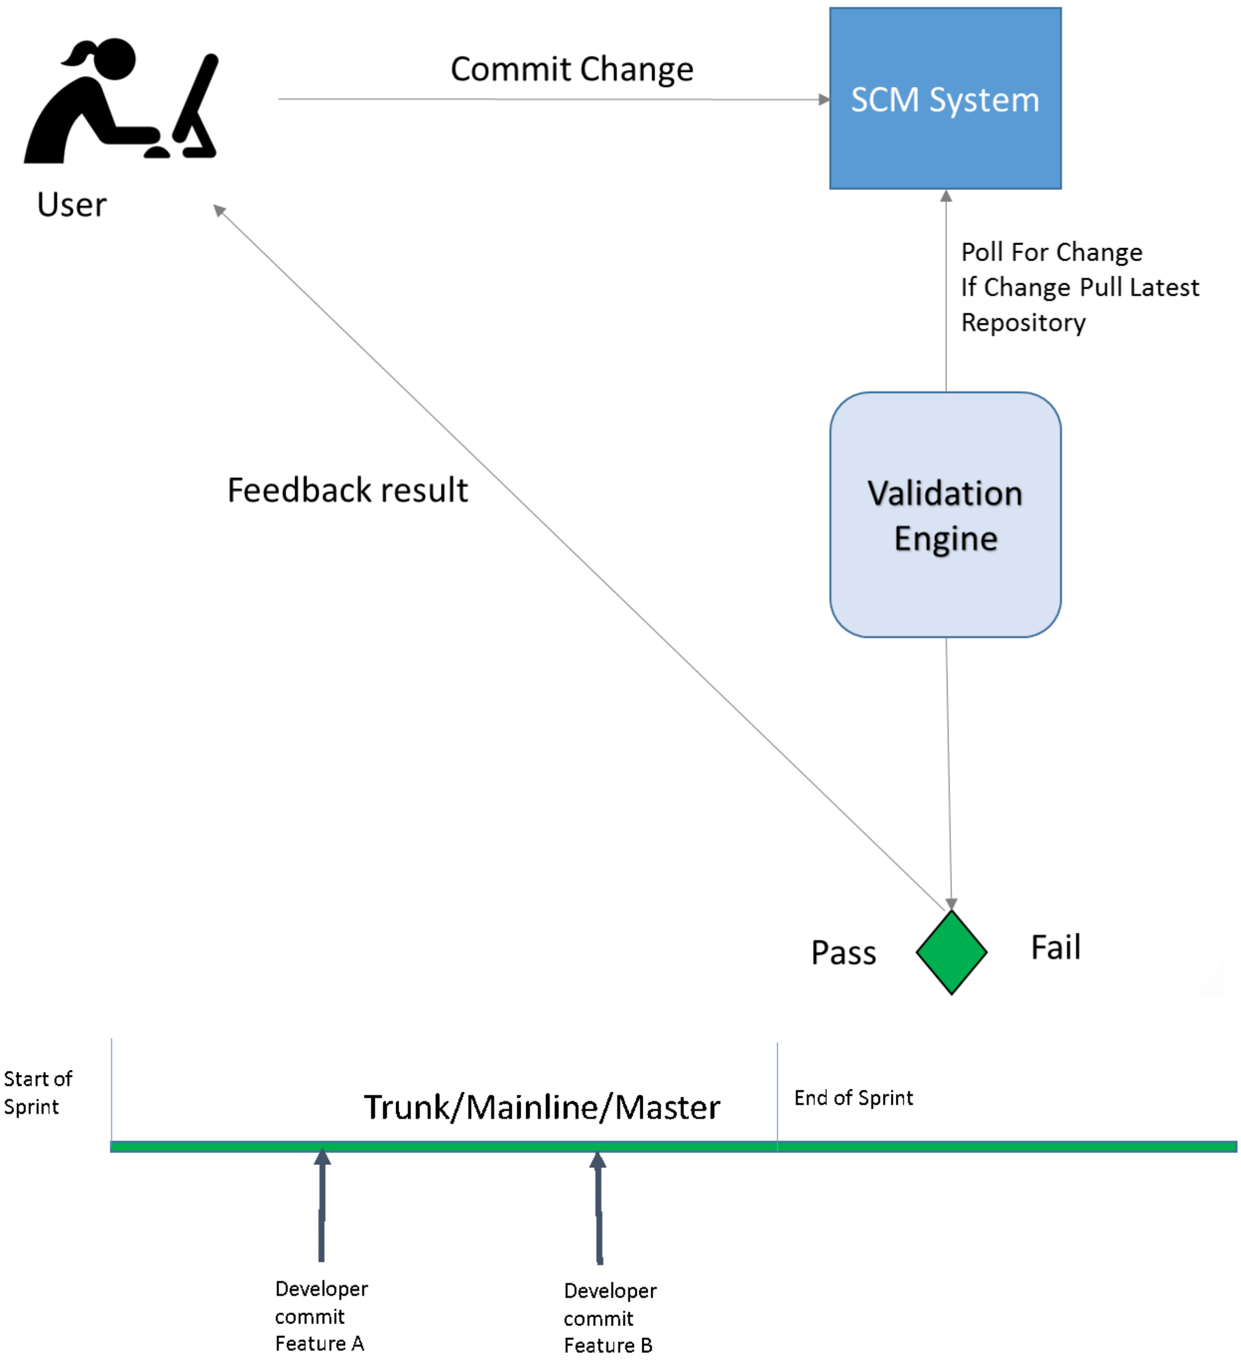
User
Commit Change
SCM System
Poll For Change
If Change Pull Latest
Repository
Validation
Engine
Feedback result
Pass
Fail
Start of
Sprint
Trunk/Mainline/Master
End of Sprint
Developer
commit
Feature A
Developer
commit
Feature B

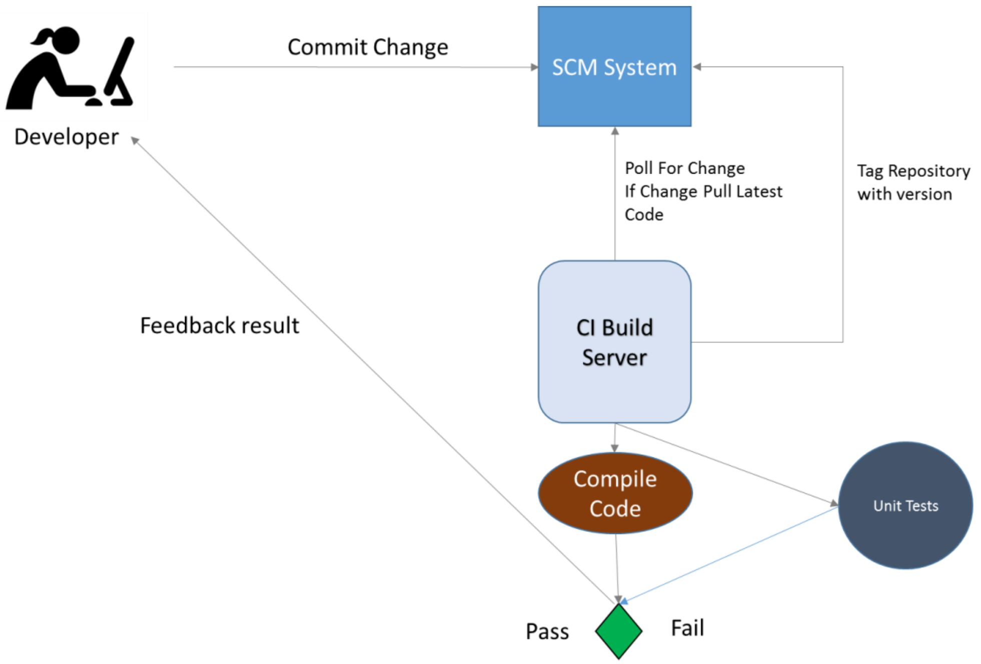
Commit Change
SCM System
Developer
Poll For Change
If Change Pull Latest
Code
Tag Repository
with version
Feedback result
CI Build
Server
Compile
Code
Unit Tests
Pass
Fail

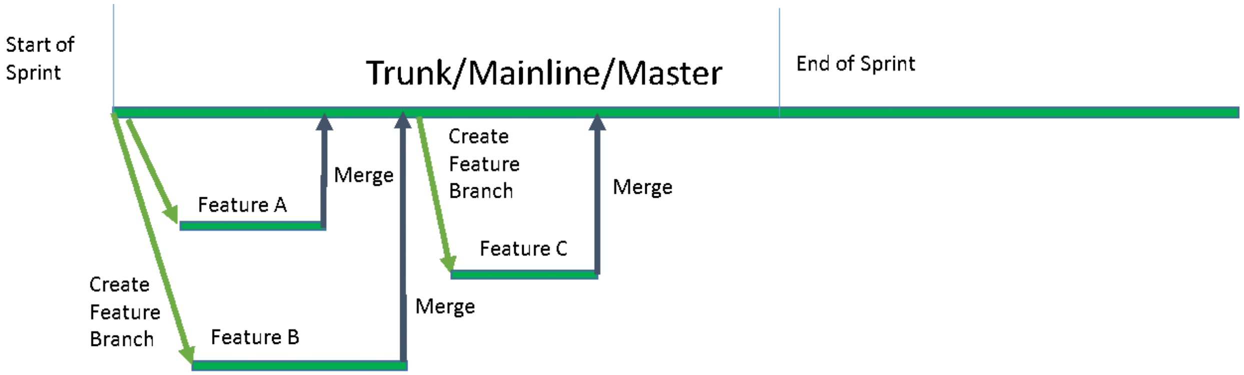
Start of
Sprint
Trunk/Mainline/Master
End of Sprint
Create
Feature
Branch
Merge
Feature A
Merge
Feature C
Create
Feature
Branch
Merge
Feature B

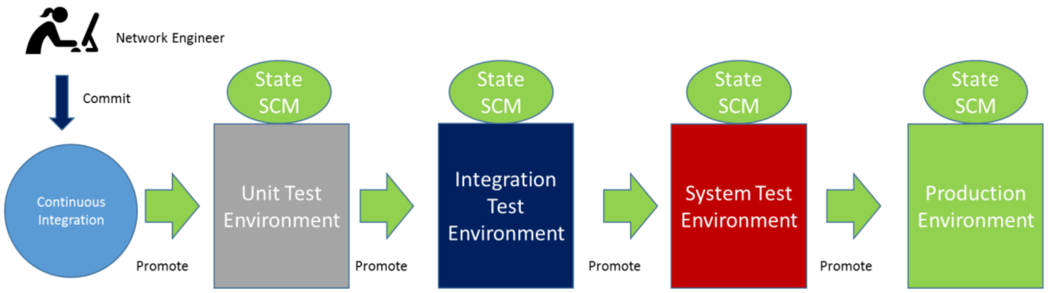
Network Engineer
Commit
Continuous
Integration
Promote
State
SCM
Unit Test
Environment
Promote
State
SCM
Integration
Test
Environment
Promote
State
SCM
System Test
Environment
Promote
State
SCM
Production
Environment

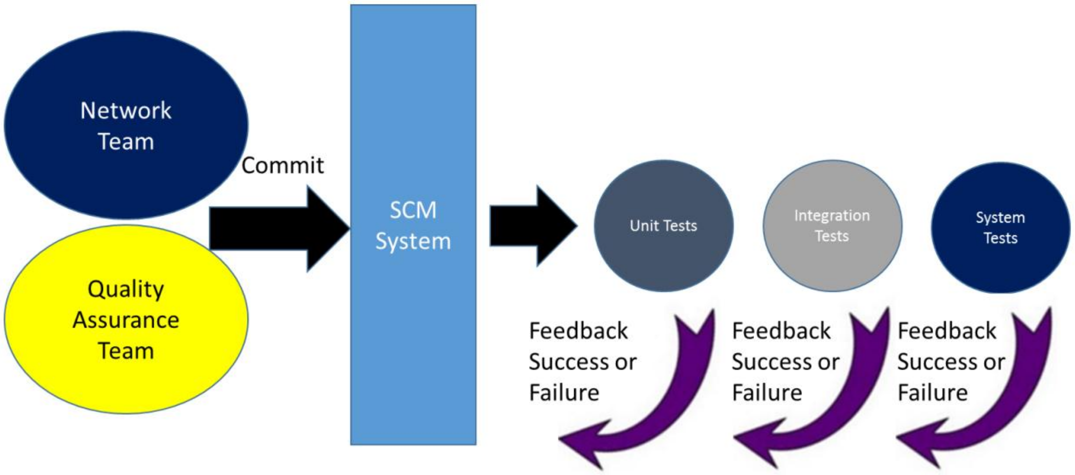
Network Team
Quality Assurance Team
Commit
SCM System
Unit Tests
Integration Tests
System Tests
Feedback Success or Failure
Feedback Success or Failure
Feedback Success or Failure

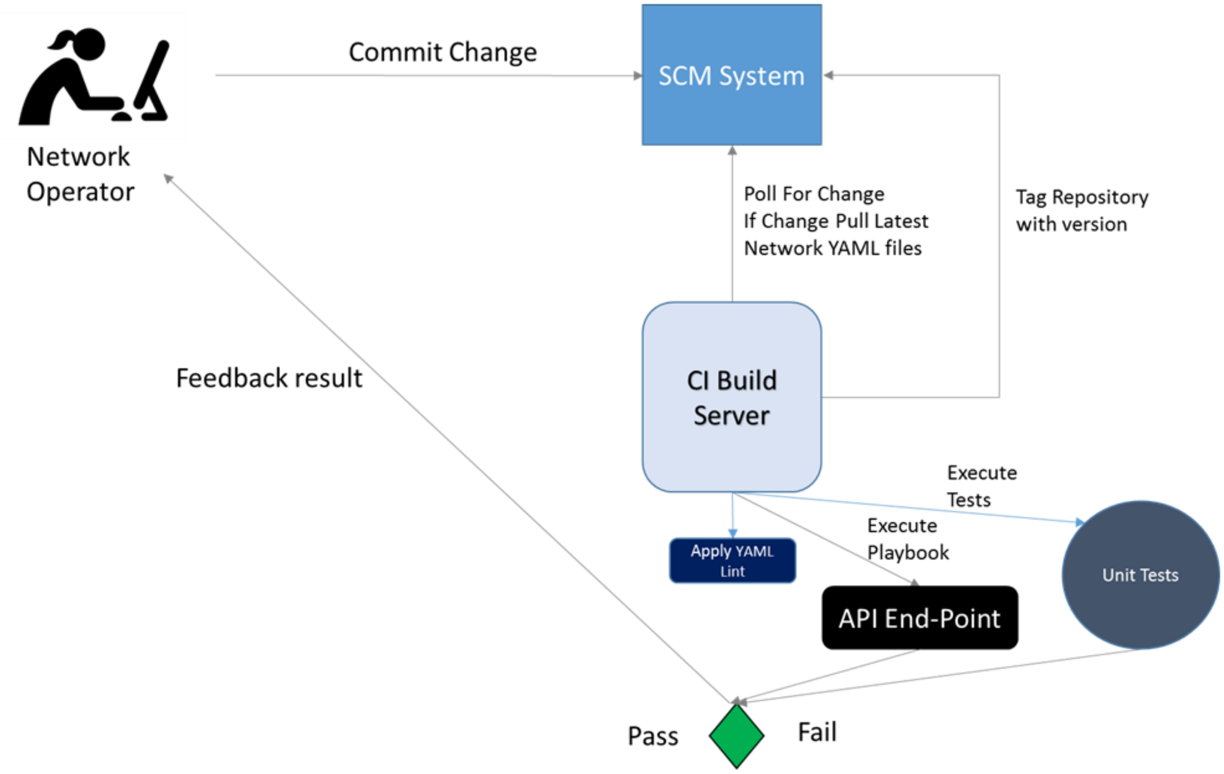
Network Operator
Commit Change
SCM System
Poll For Change
If Change Pull Latest
Network YAML files
Tag Repository
with version
CI Build Server
Feedback result
Execute
Tests
Execute
Playbook
Apply YAML Lint
API End-Point
Unit Tests
Pass
Fail

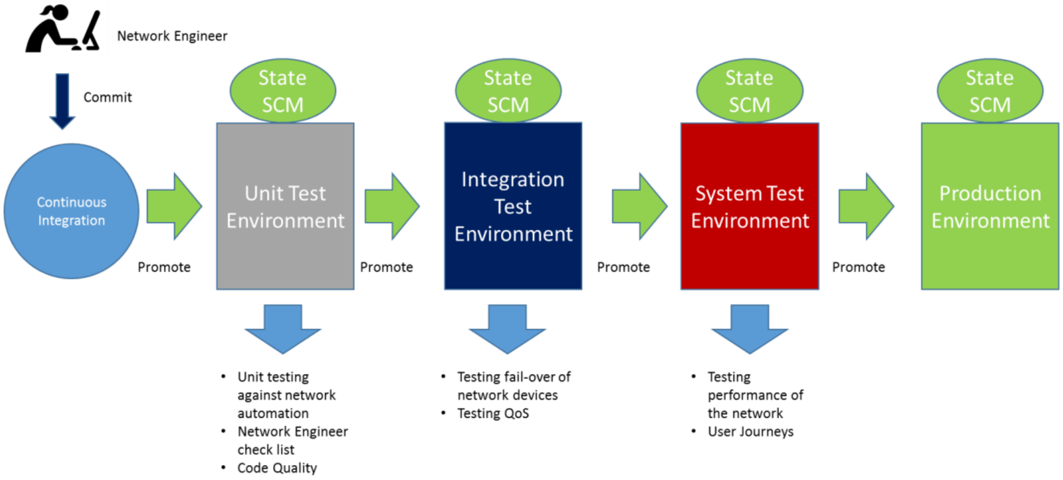
Network Engineer
Commit
Continuous Integration
Promote
State SCM
Unit Test Environment
Promote
State SCM
Integration Test Environment
Promote
State SCM
System Test Environment
Promote
State SCM
Production Environment
• Unit testing against network automation
• Network Engineer check list
• Code Quality
• Testing fail-over of network devices
• Testing QoS
• Testing performance of the network
• User Journeys

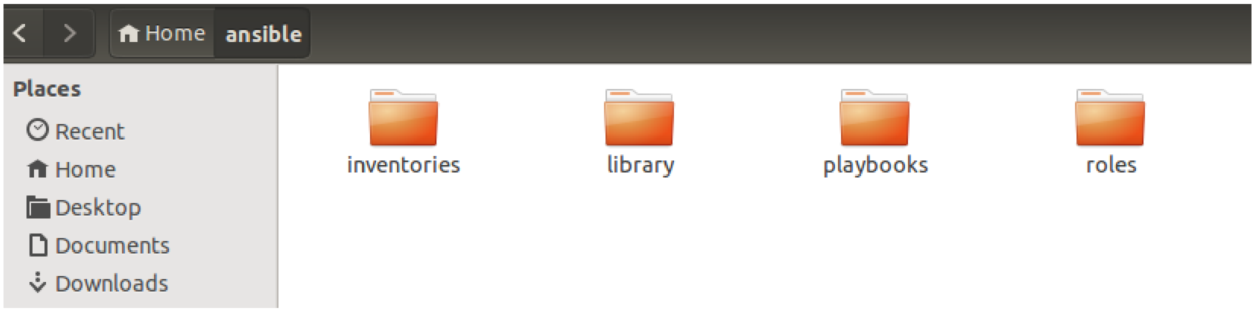
Home
ansible
Places
Recent
Home
Desktop
Documents
Downloads
inventories
library
playbooks
roles

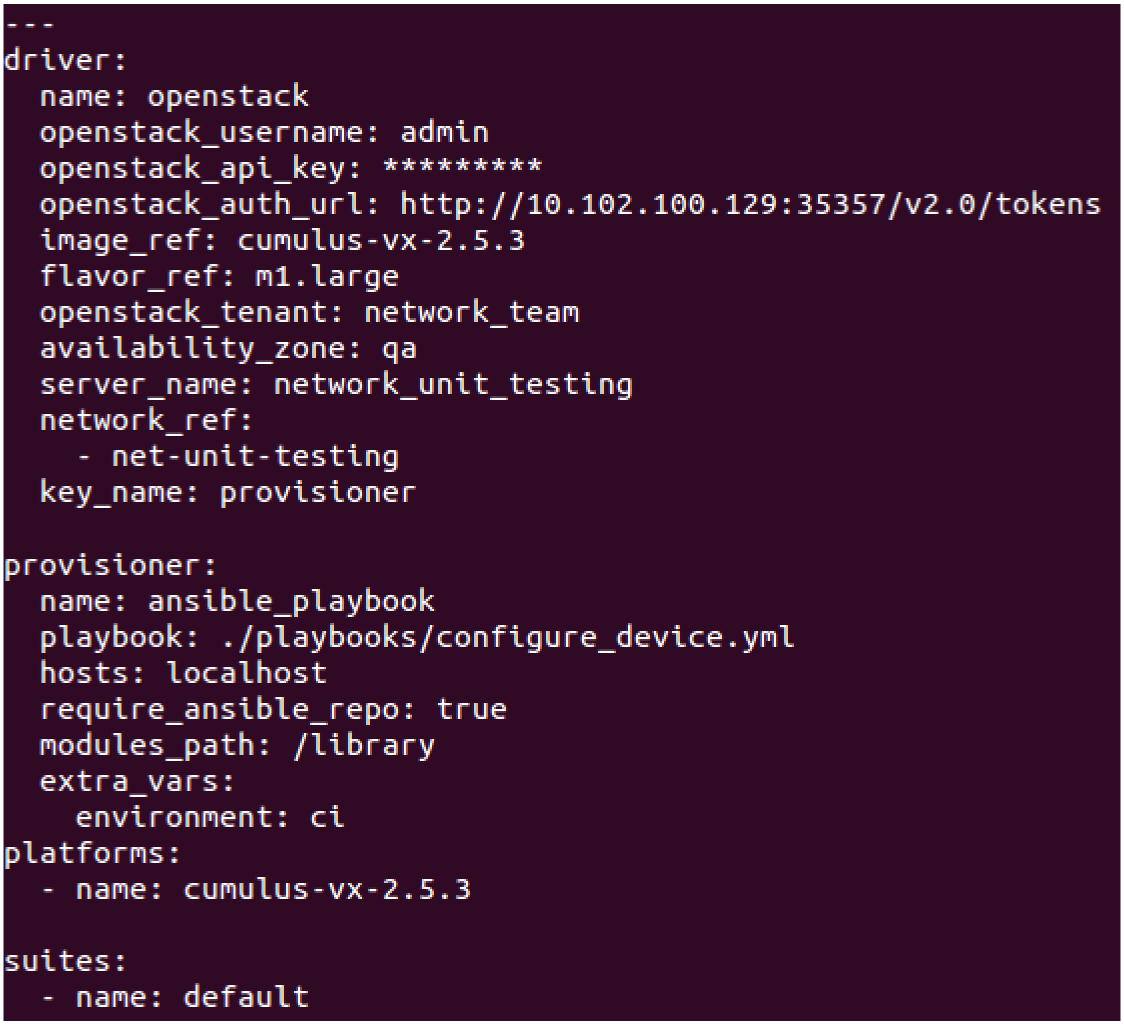

```
---
driver:
  name: openstack
  openstack_username: admin
  openstack_api_key: *********
  openstack_auth_url: http://10.102.100.129:35357/v2.0/tokens
  image_ref: cumulus-vx-2.5.3
  flavor_ref: m1.large
  openstack_tenant: network_team
  availability_zone: qa
  server_name: network_unit_testing
  network_ref:
    - net-unit-testing
  key_name: provisioner

provisioner:
  name: ansible_playbook
  playbook: ./playbooks/configure_device.yml
  hosts: localhost
  require_ansible_repo: true
  modules_path: /library
  extra_vars:
    environment: ci
platforms:
  - name: cumulus-vx-2.5.3

suites:
  - name: default
```

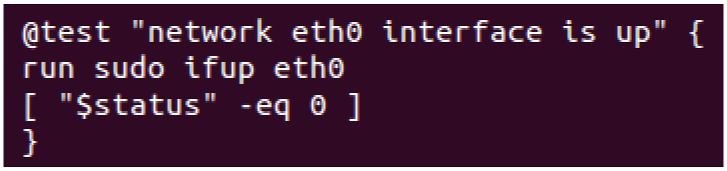

```
@test "network eth0 interface is up" {
run sudo ifup eth0
[ "$status" -eq 0 ]
}
```

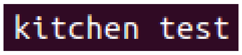

```
kitchen test
```

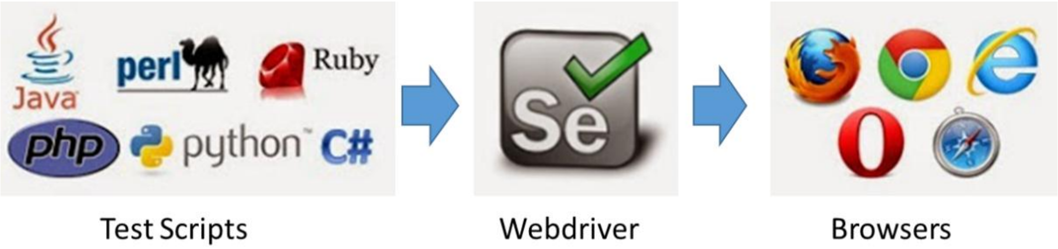

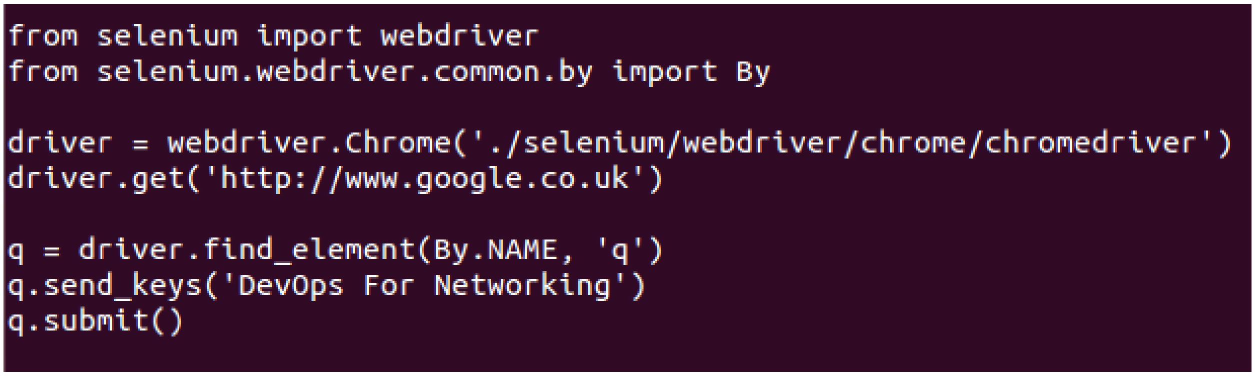

```
from selenium import webdriver
from selenium.webdriver.common.by import By

driver = webdriver.Chrome('./selenium/webdriver/chrome/chromedriver')
driver.get('http://www.google.co.uk')

q = driver.find_element(By.NAME, 'q')
q.send_keys('DevOps For Networking')
q.submit()
```

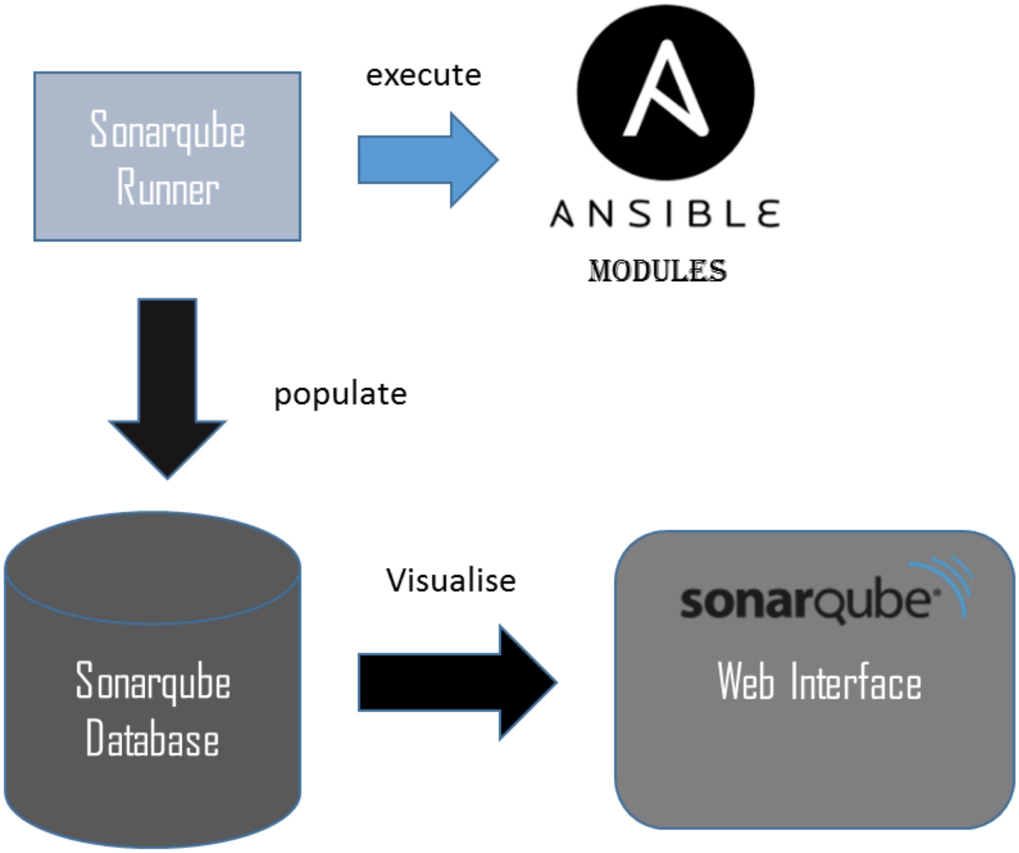

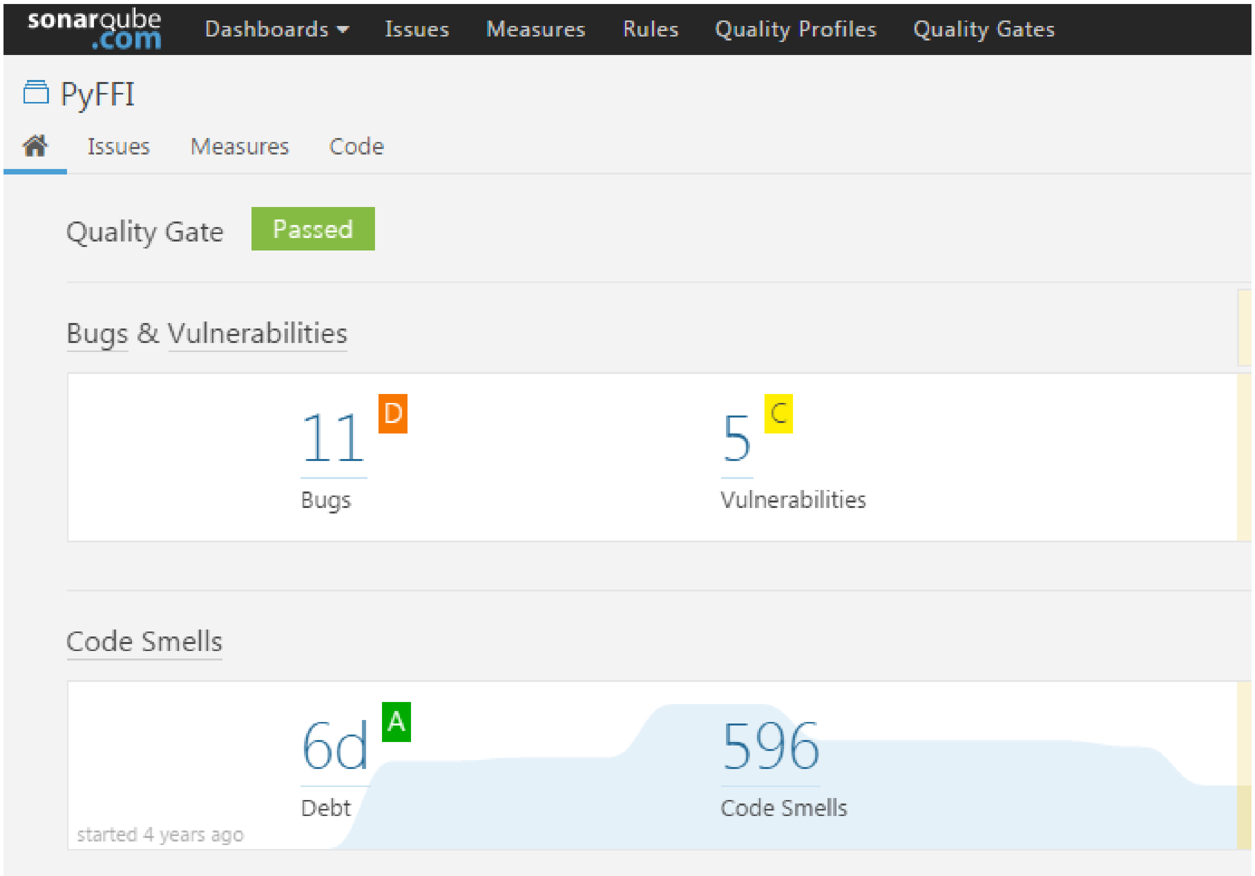

Quality Gate Passed

Bugs & Vulnerabilities

11 D Bugs

5 C Vulnerabilities

Code Smells

6d A Debt

596 Code Smells

started 4 years ago

Chapter 9: Using Continuous Delivery Pipelines to Deploy Network Changes

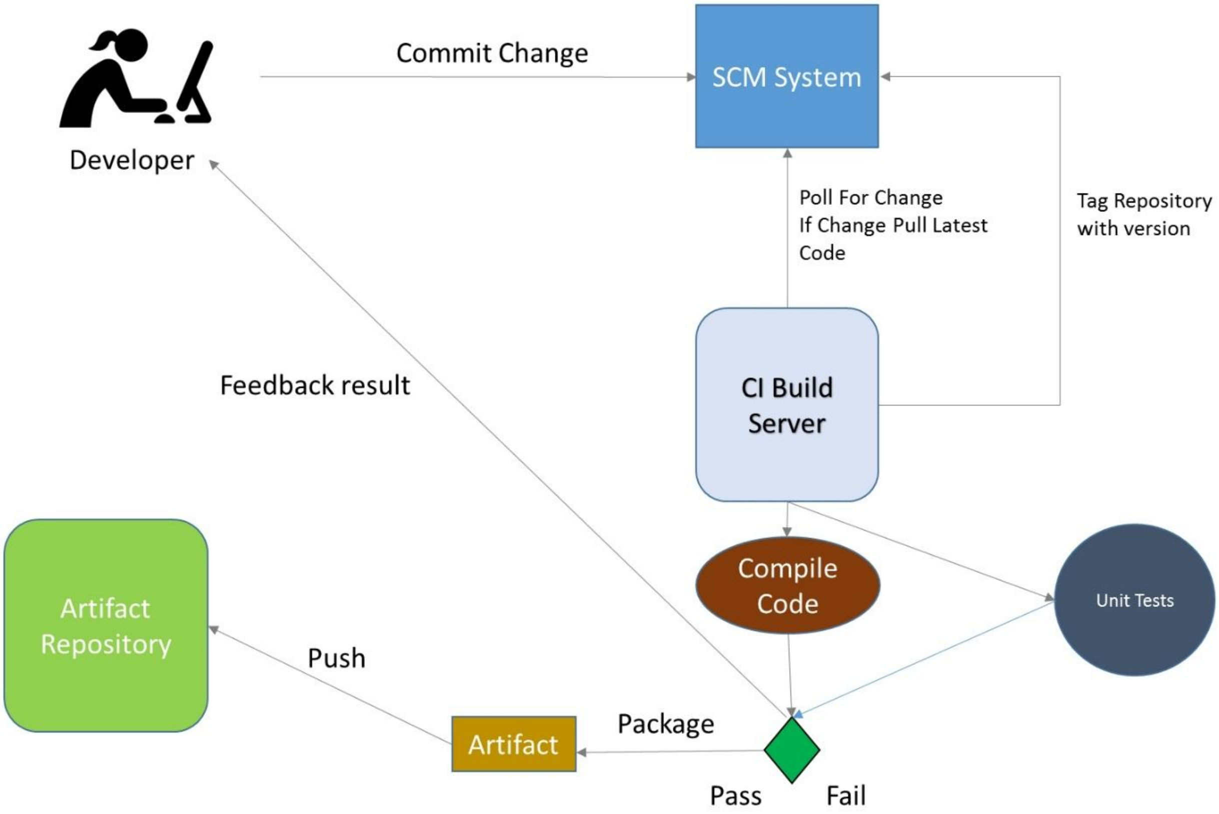

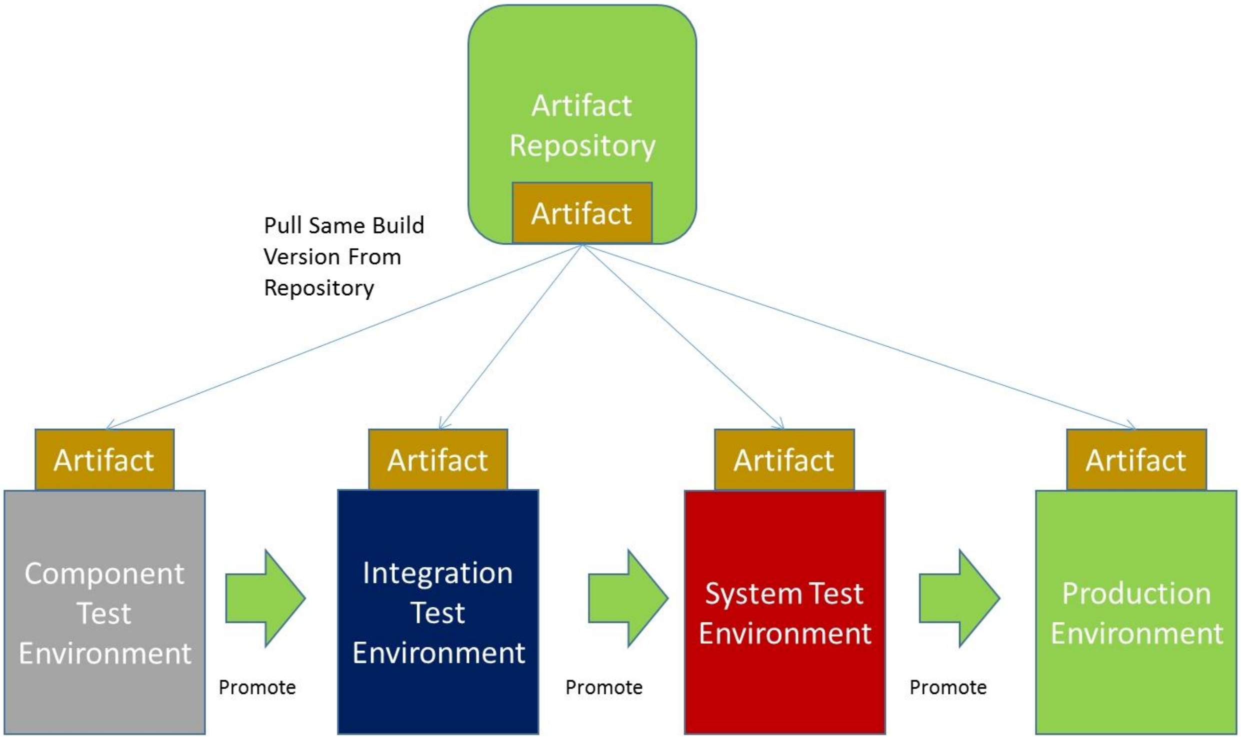

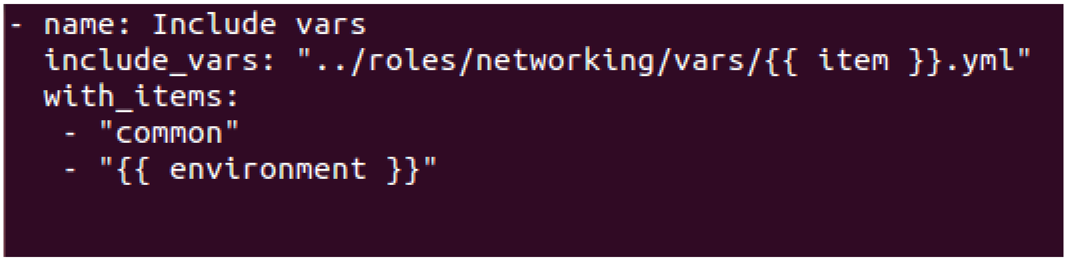

```
- name: Include vars
  include_vars: "../roles/networking/vars/{{ item }}.yml"
  with_items:
   - "common"
   - "{{ environment }}"
```

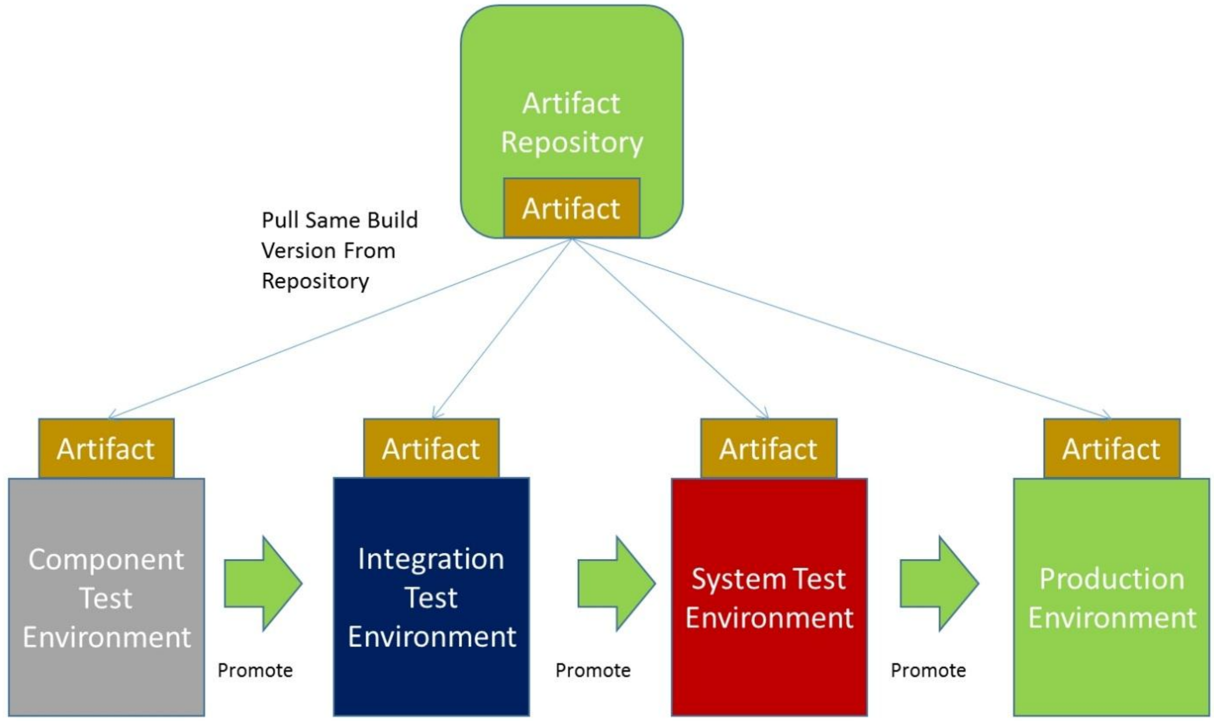

```
ansible-playbook
configure_env.yml
-e environment=comp
```

```
ansible-playbook
configure_env.yml
-e environment=int
```

```
ansible-playbook
configure_env.yml
-e environment=sys
```

```
ansible-playbook
configure_env.yml
-e environment=prod
```

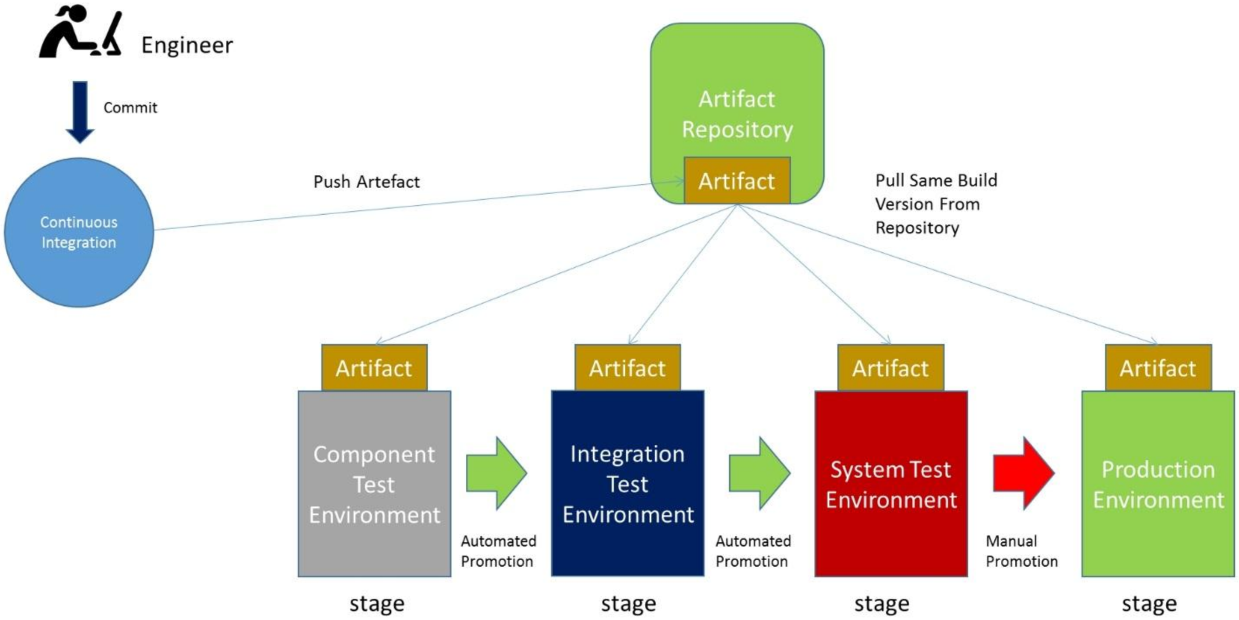

Continuous Delivery

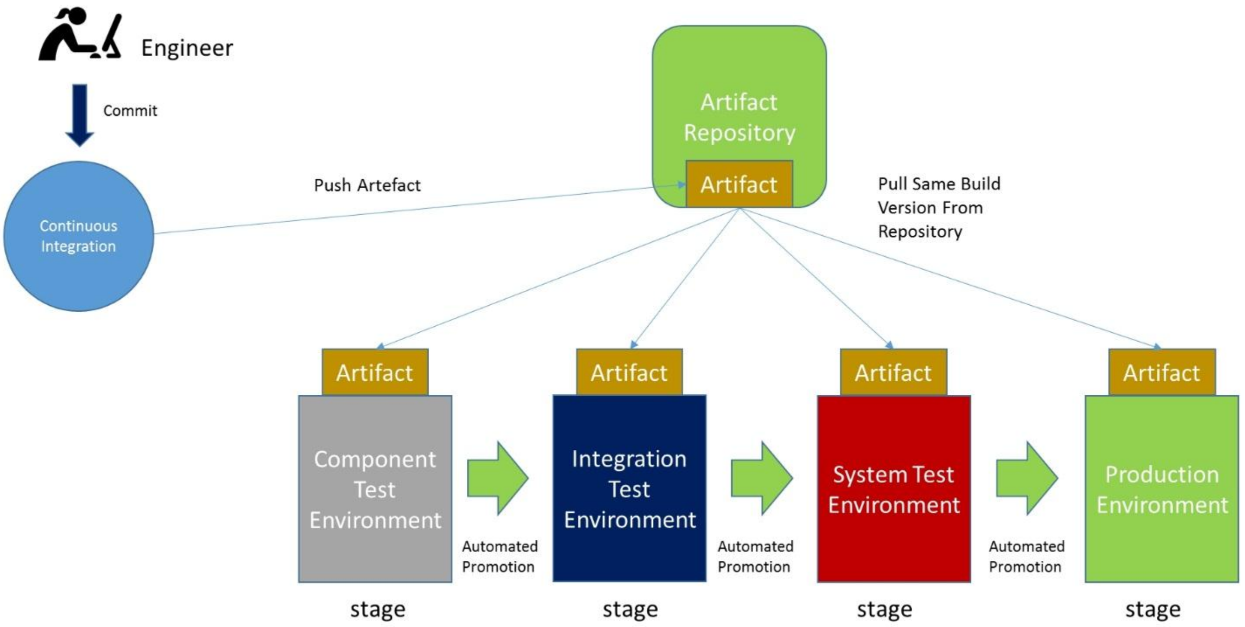

Continuous Deployment

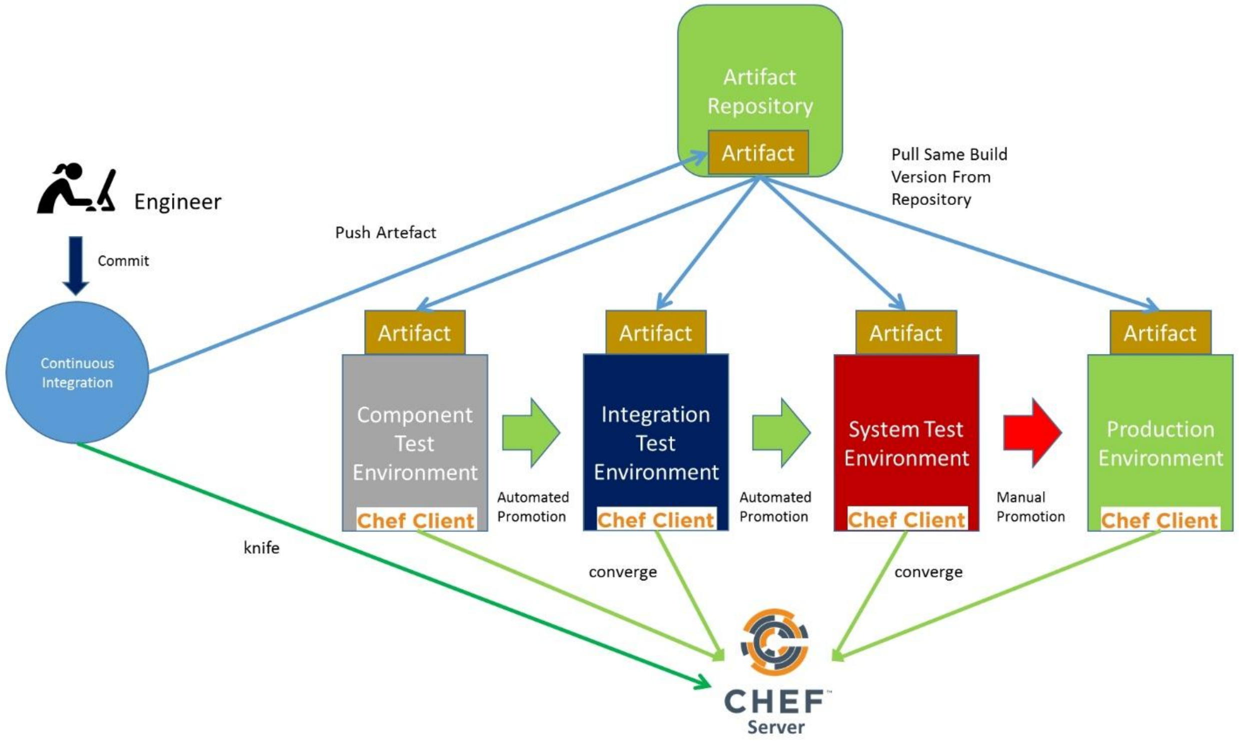

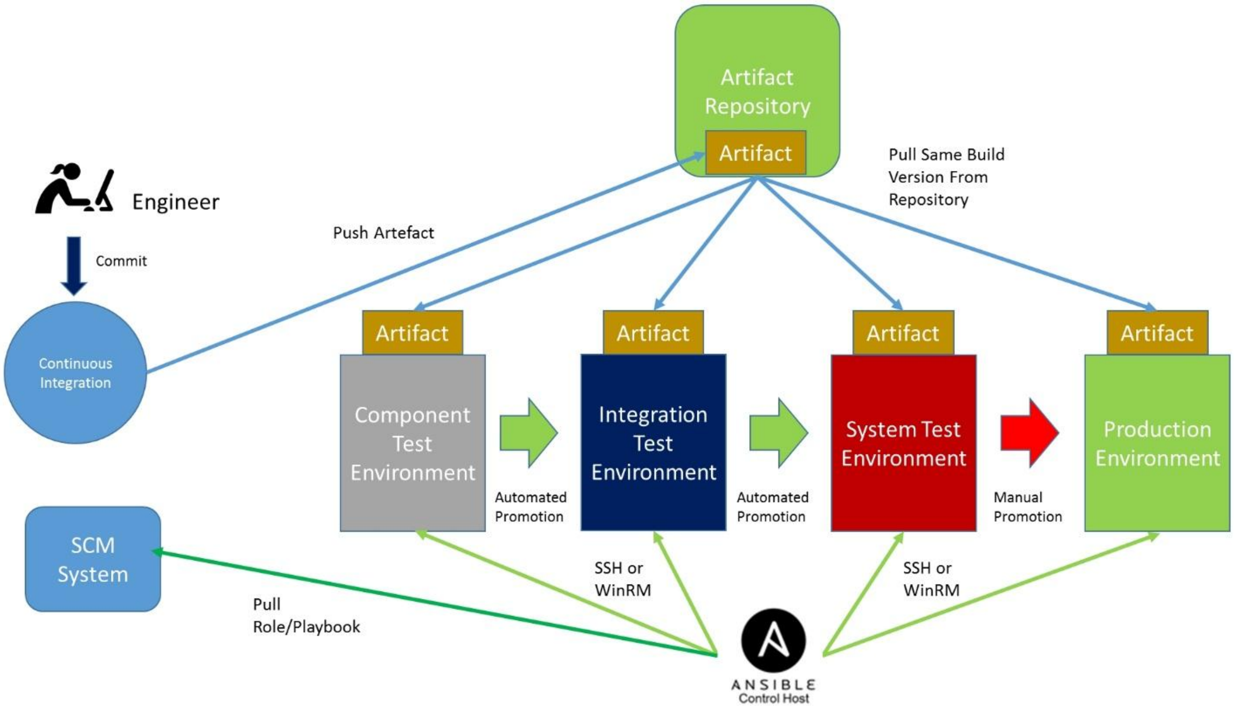

Artifact Repository
Artifact
Pull Same Build Version From Repository
Engineer
Push Artefact
Commit
Continuous Integration
Artifact
Artifact
Artifact
Artifact
Component Test Environment
Integration Test Environment
System Test Environment
Production Environment
Automated Promotion
Automated Promotion
Manual Promotion
SCM System
SSH or WinRM
SSH or WinRM
Pull Role/Playbook
ANSIBLE
Control Host

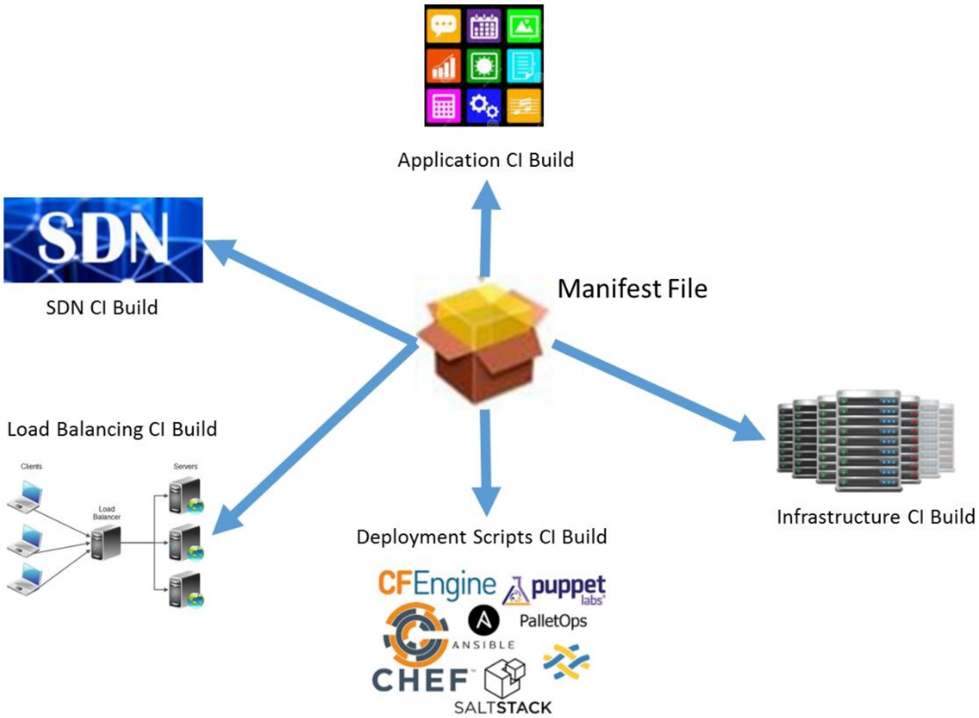

Application CI Build
SDN
SDN CI Build
Manifest File
Load Balancing CI Build
Clients
Load Balancer
Servers
Deployment Scripts CI Build
Infrastructure CI Build
CFEngine
puppet labs
ANSIBLE
PalletOps
CHEF
SALTSTACK

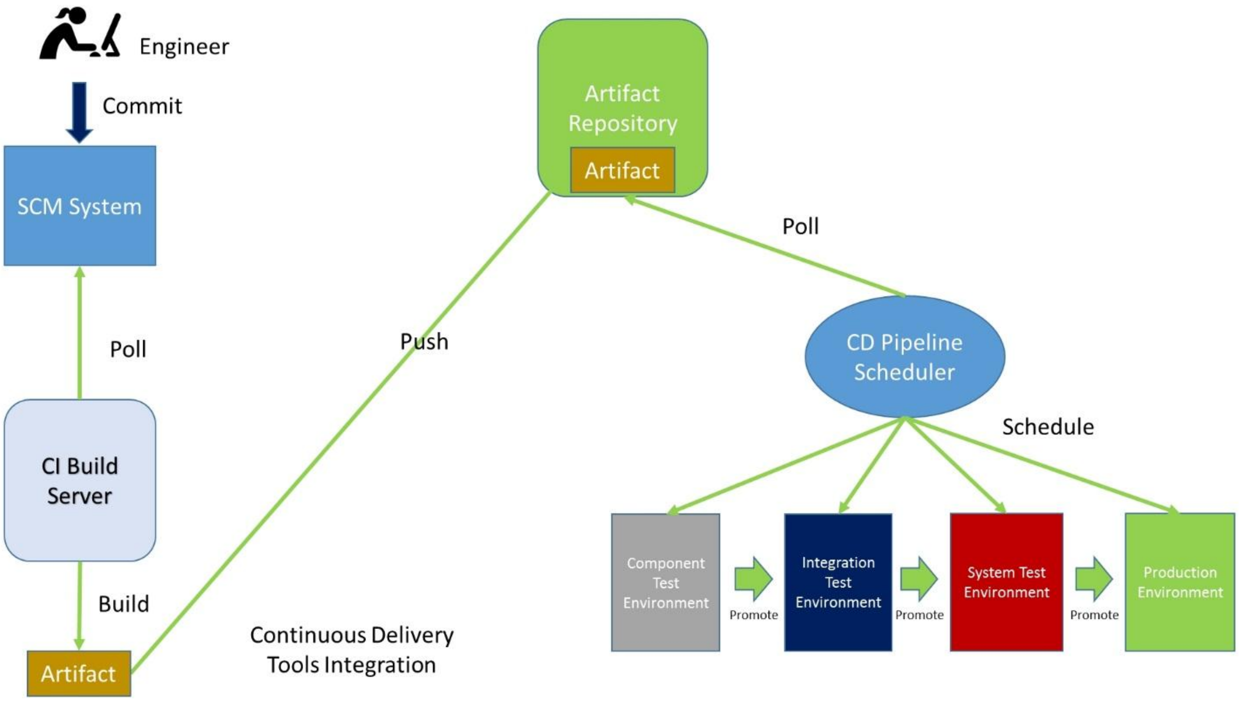
Engineer
Commit
SCM System
Poll
CI Build Server
Build
Artifact
Push
Artifact Repository
Artifact
Poll
CD Pipeline Scheduler
Schedule
Component Test Environment
Promote
Integration Test Environment
Promote
System Test Environment
Promote
Production Environment
Continuous Delivery Tools Integration

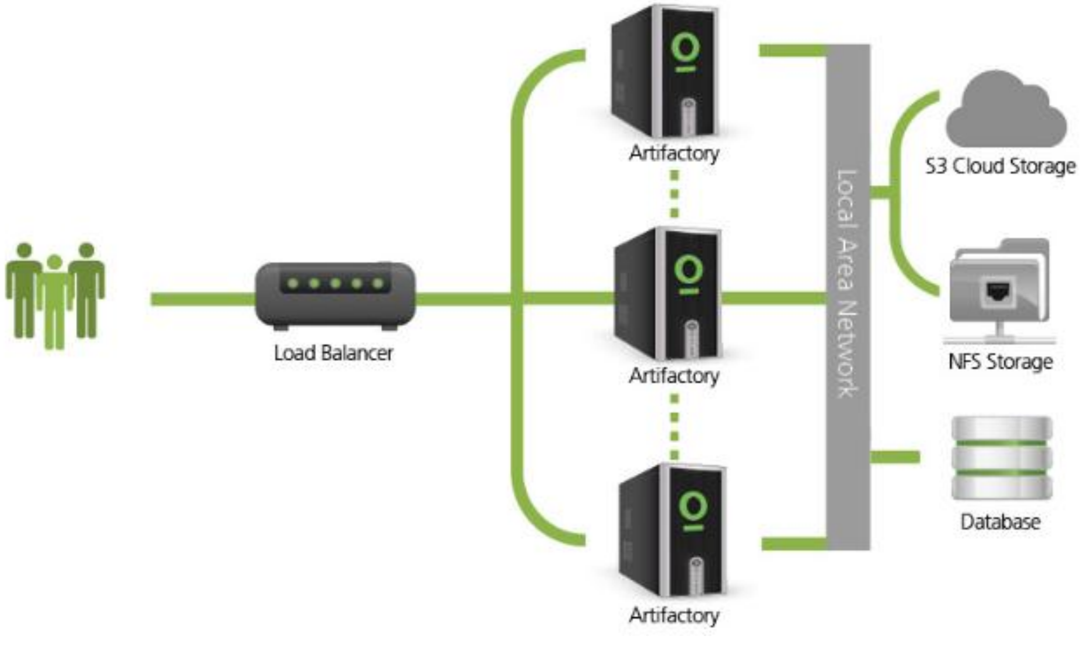
Load Balancer
Artifactory
Artifactory
Artifactory
Local Area Network
S3 Cloud Storage
NFS Storage
Database

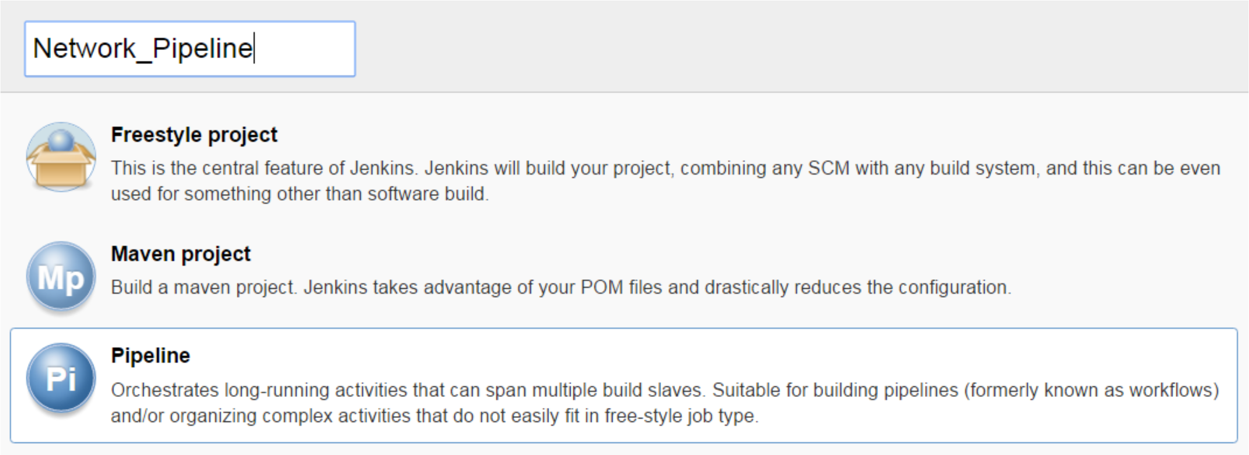
Network_Pipeline

Freestyle project

This is the central feature of Jenkins. Jenkins will build your project, combining any SCM with any build system, and this can be even used for something other than software build.

Maven project

Build a maven project. Jenkins takes advantage of your POM files and drastically reduces the configuration.

Pipeline

Orchestrates long-running activities that can span multiple build slaves. Suitable for building pipelines (formerly known as workflows) and/or organizing complex activities that do not easily fit in free-style job type.

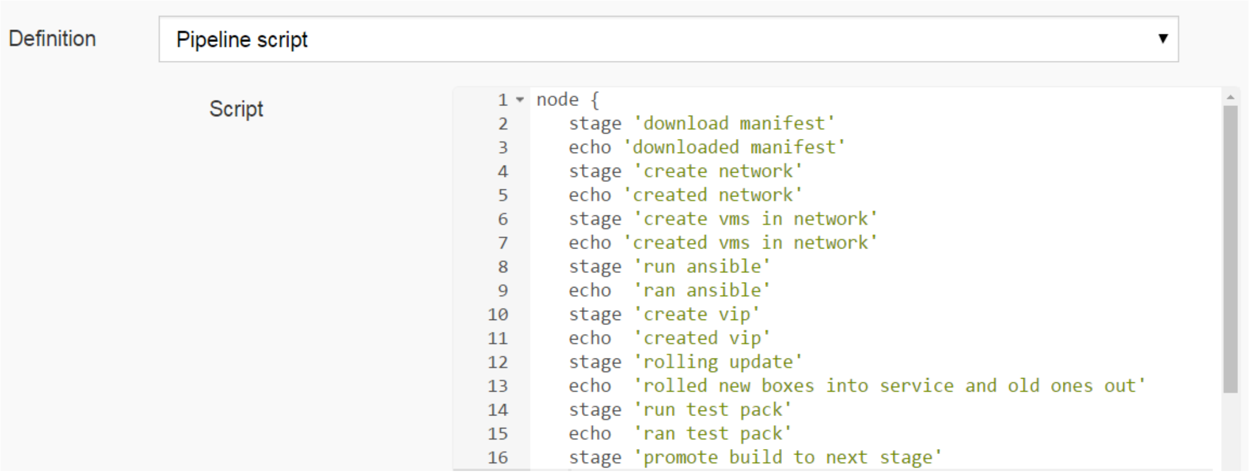
Definition: Pipeline script

Script

```
node {
   stage 'download manifest'
   echo 'downloaded manifest'
   stage 'create network'
   echo 'created network'
   stage 'create vms in network'
   echo 'created vms in network'
   stage 'run ansible'
   echo  'ran ansible'
   stage 'create vip'
   echo  'created vip'
   stage 'rolling update'
   echo  'rolled new boxes into service and old ones out'
   stage 'run test pack'
   echo  'ran test pack'
   stage 'promote build to next stage'
```

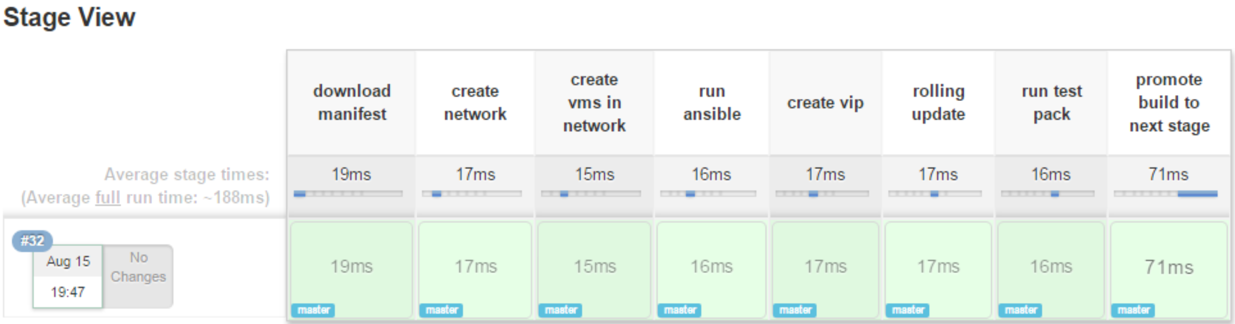
Stage View

	download manifest	create network	create vms in network	run ansible	create vip	rolling update	run test pack	promote build to next stage
Average stage times: (Average full run time: ~188ms)	19ms	17ms	15ms	16ms	17ms	17ms	16ms	71ms
#32 Aug 15 19:47 No Changes	19ms master	17ms master	15ms master	16ms master	17ms master	17ms master	16ms master	71ms master

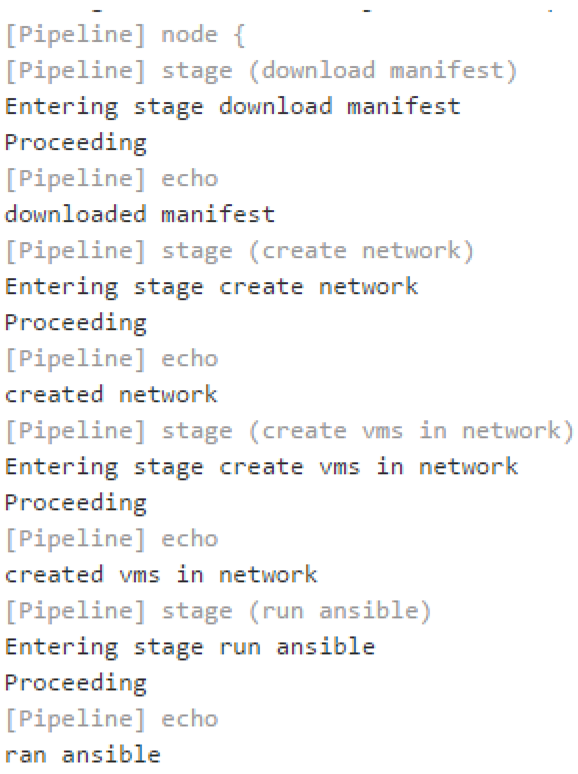

```
[Pipeline] node {
[Pipeline] stage (download manifest)
Entering stage download manifest
Proceeding
[Pipeline] echo
downloaded manifest
[Pipeline] stage (create network)
Entering stage create network
Proceeding
[Pipeline] echo
created network
[Pipeline] stage (create vms in network)
Entering stage create vms in network
Proceeding
[Pipeline] echo
created vms in network
[Pipeline] stage (run ansible)
Entering stage run ansible
Proceeding
[Pipeline] echo
ran ansible
```

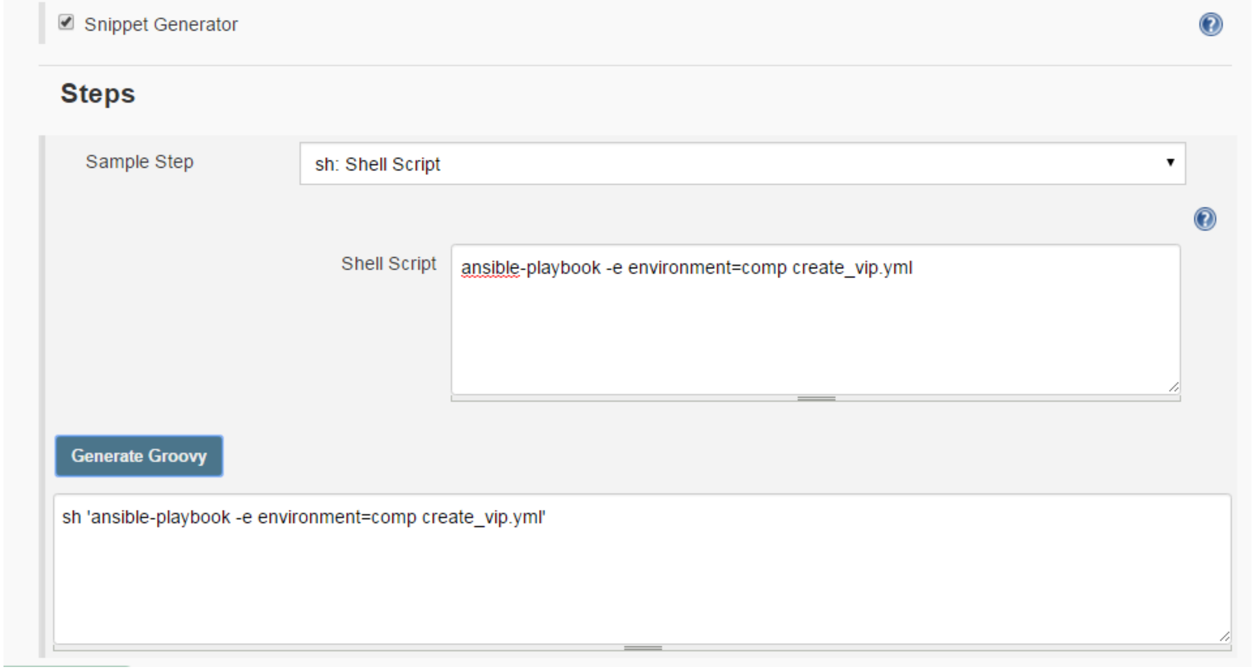

☑ Snippet Generator

Sample Step: sh: Shell Script

Shell Script: ansible-playbook -e environment=comp create_vip.yml

```
sh 'ansible-playbook -e environment=comp create_vip.yml'
```

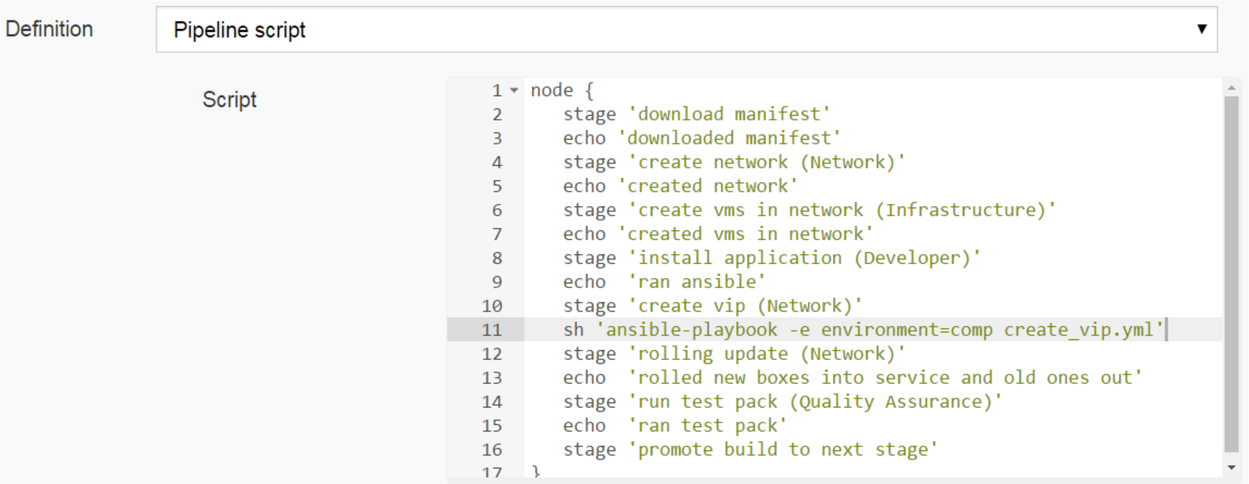

Definition: Pipeline script

Script

```
node {
   stage 'download manifest'
   echo 'downloaded manifest'
   stage 'create network (Network)'
   echo 'created network'
   stage 'create vms in network (Infrastructure)'
   echo 'created vms in network'
   stage 'install application (Developer)'
   echo  'ran ansible'
   stage 'create vip (Network)'
   sh 'ansible-playbook -e environment=comp create_vip.yml'
   stage 'rolling update (Network)'
   echo  'rolled new boxes into service and old ones out'
   stage 'run test pack (Quality Assurance)'
   echo  'ran test pack'
   stage 'promote build to next stage'
}
```

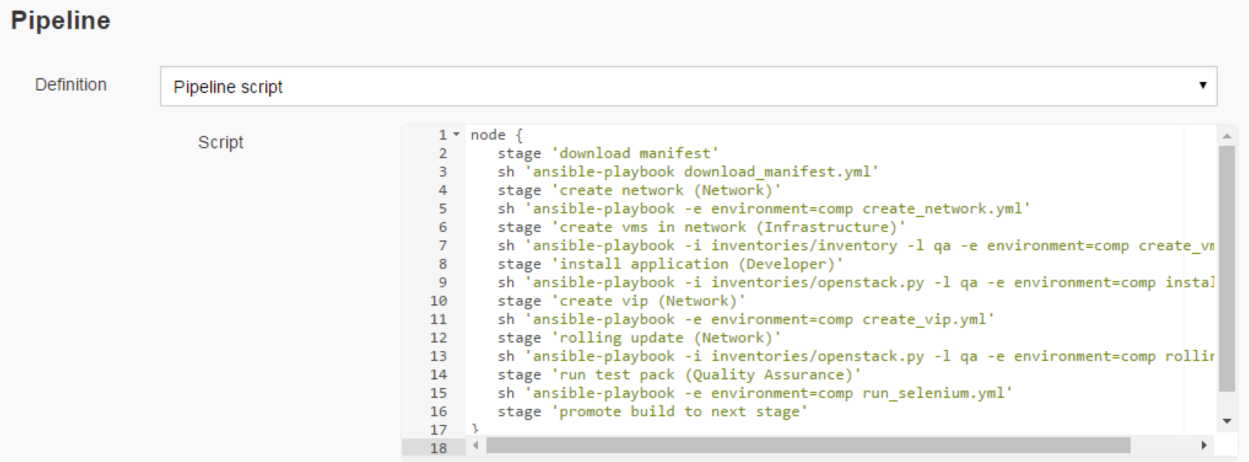

Pipeline

Definition: Pipeline script

Script

```
node {
   stage 'download manifest'
   sh 'ansible-playbook download_manifest.yml'
   stage 'create network (Network)'
   sh 'ansible-playbook -e environment=comp create_network.yml'
   stage 'create vms in network (Infrastructure)'
   sh 'ansible-playbook -i inventories/inventory -l qa -e environment=comp create_vm
   stage 'install application (Developer)'
   sh 'ansible-playbook -i inventories/openstack.py -l qa -e environment=comp instal
   stage 'create vip (Network)'
   sh 'ansible-playbook -e environment=comp create_vip.yml'
   stage 'rolling update (Network)'
   sh 'ansible-playbook -i inventories/openstack.py -l qa -e environment=comp rollir
   stage 'run test pack (Quality Assurance)'
   sh 'ansible-playbook -e environment=comp run_selenium.yml'
   stage 'promote build to next stage'
}
```

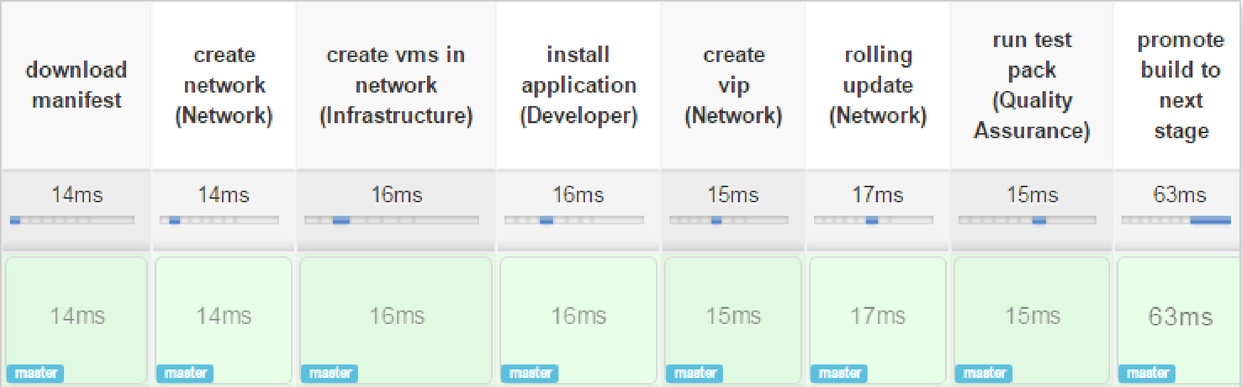

download manifest	create network (Network)	create vms in network (Infrastructure)	install application (Developer)	create vip (Network)	rolling update (Network)	run test pack (Quality Assurance)	promote build to next stage
14ms	14ms	16ms	16ms	15ms	17ms	15ms	63ms
14ms master	14ms master	16ms master	16ms master	15ms master	17ms master	15ms master	63ms master

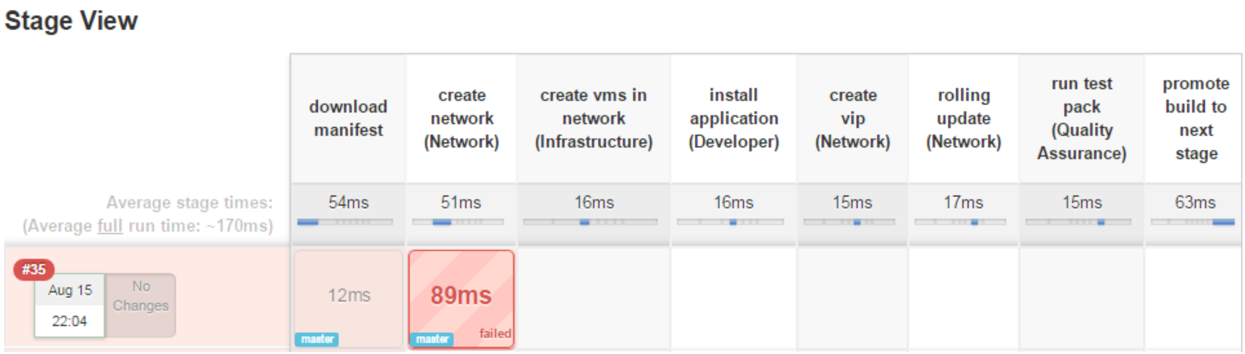

Stage View

	download manifest	create network (Network)	create vms in network (Infrastructure)	install application (Developer)	create vip (Network)	rolling update (Network)	run test pack (Quality Assurance)	promote build to next stage
Average stage times: (Average full run time: ~170ms)	54ms	51ms	16ms	16ms	15ms	17ms	15ms	63ms
#35 Aug 15 22:04 No Changes	12ms master	89ms master failed						

Chapter 10: The Impact Of Containers On Networking

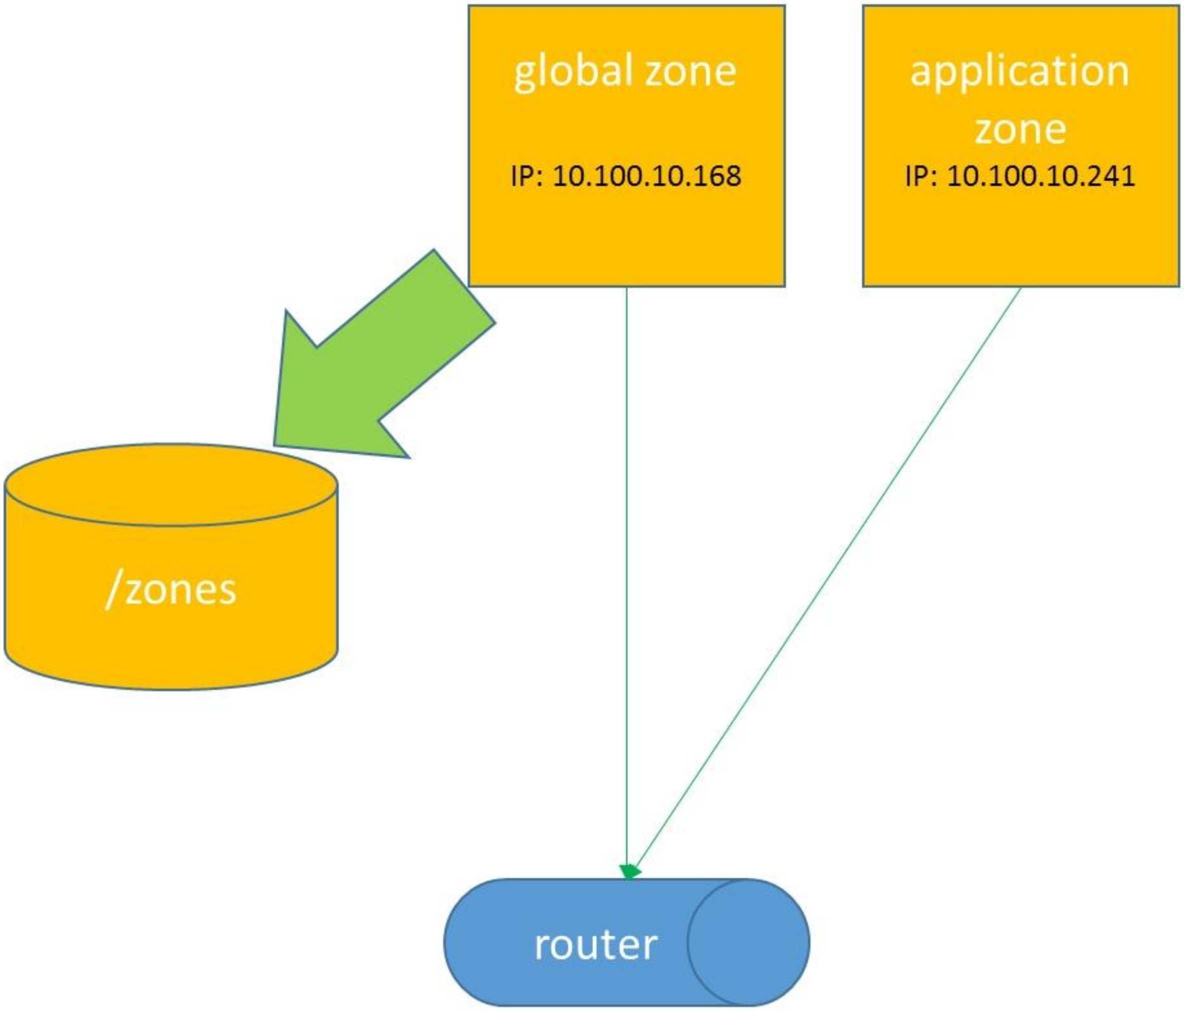

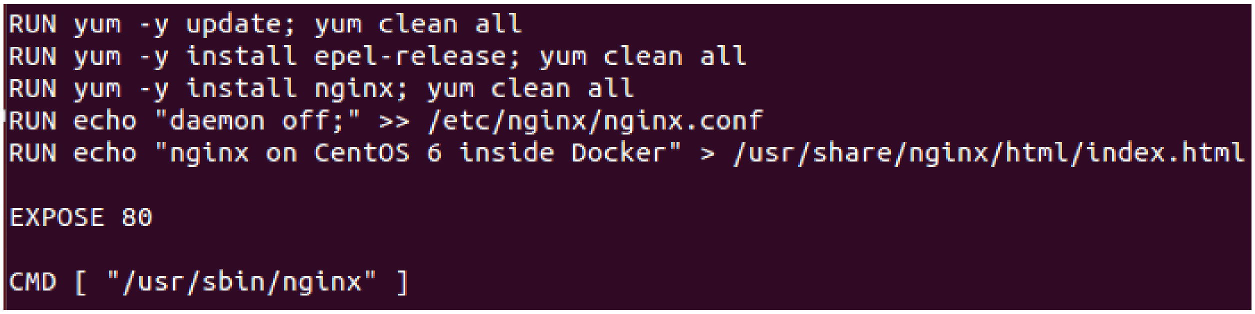

```
RUN yum -y update; yum clean all
RUN yum -y install epel-release; yum clean all
RUN yum -y install nginx; yum clean all
RUN echo "daemon off;" >> /etc/nginx/nginx.conf
RUN echo "nginx on CentOS 6 inside Docker" > /usr/share/nginx/html/index.html

EXPOSE 80

CMD [ "/usr/sbin/nginx" ]
```

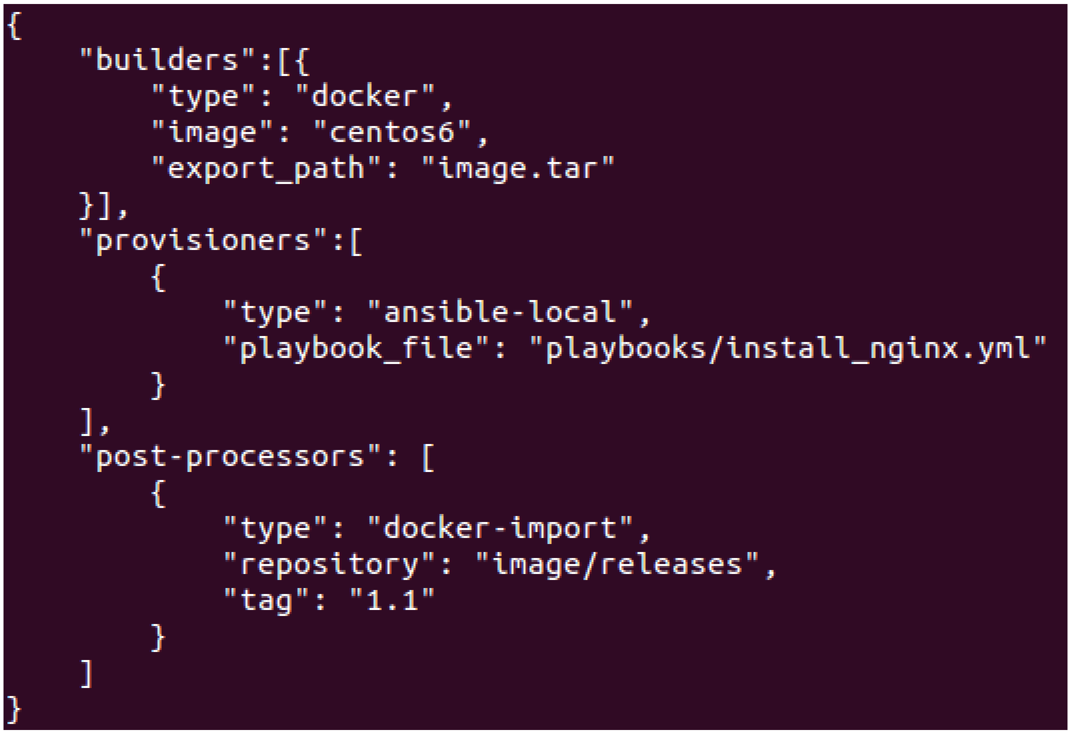

```
{
    "builders":[{
        "type": "docker",
        "image": "centos6",
        "export_path": "image.tar"
    }],
    "provisioners":[
        {
            "type": "ansible-local",
            "playbook_file": "playbooks/install_nginx.yml"
        }
    ],
    "post-processors": [
        {
            "type": "docker-import",
            "repository": "image/releases",
            "tag": "1.1"
        }
    ]
}
```

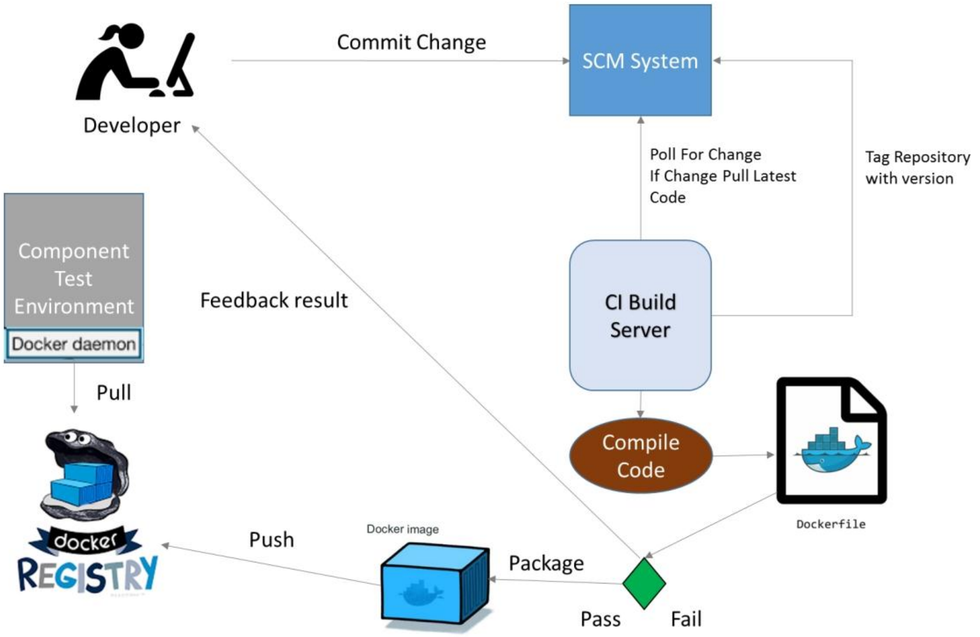

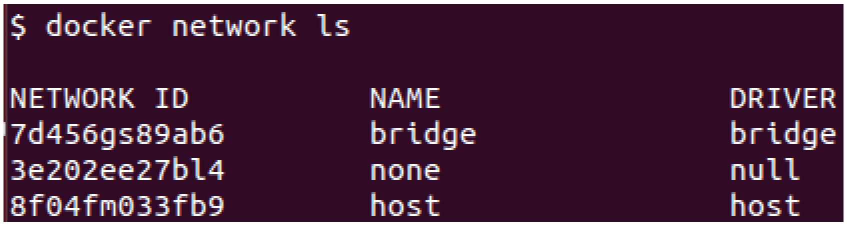

```
$ docker network ls

NETWORK ID          NAME                DRIVER
7d456gs89ab6        bridge              bridge
3e202ee27bl4        none                null
8f04fm033fb9        host                host
```

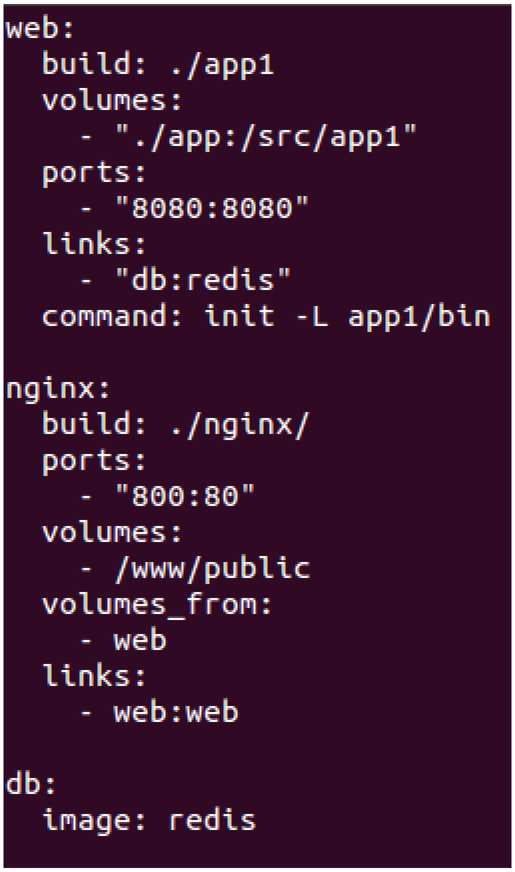

```
web:
  build: ./app1
  volumes:
    - "./app:/src/app1"
  ports:
    - "8080:8080"
  links:
    - "db:redis"
  command: init -L app1/bin

nginx:
  build: ./nginx/
  ports:
    - "800:80"
  volumes:
    - /www/public
  volumes_from:
    - web
  links:
    - web:web

db:
  image: redis
```

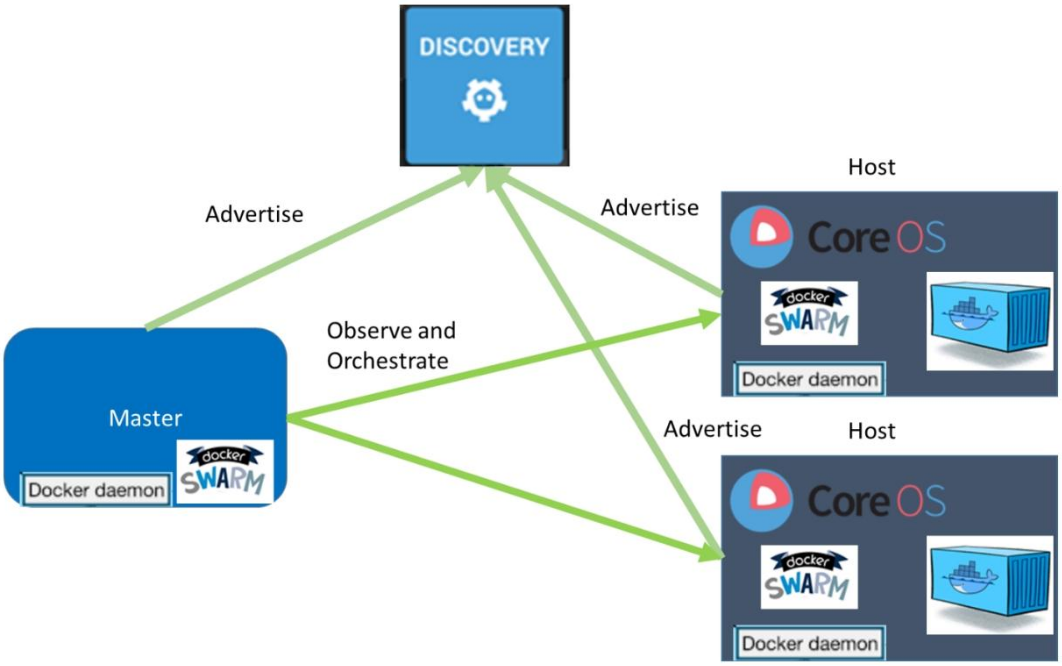

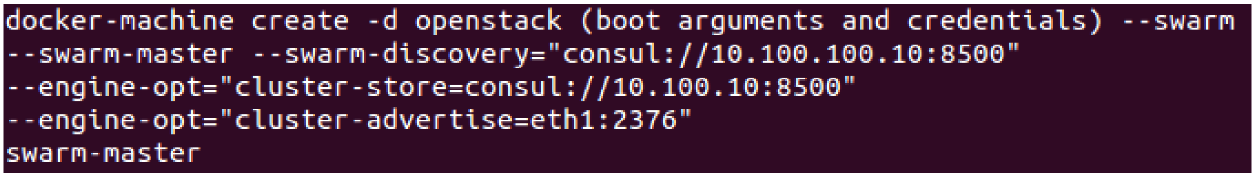

```
docker-machine create -d openstack (boot arguments and credentials) --swarm
--swarm-master --swarm-discovery="consul://10.100.100.10:8500"
--engine-opt="cluster-store=consul://10.100.10:8500"
--engine-opt="cluster-advertise=eth1:2376"
swarm-master
```

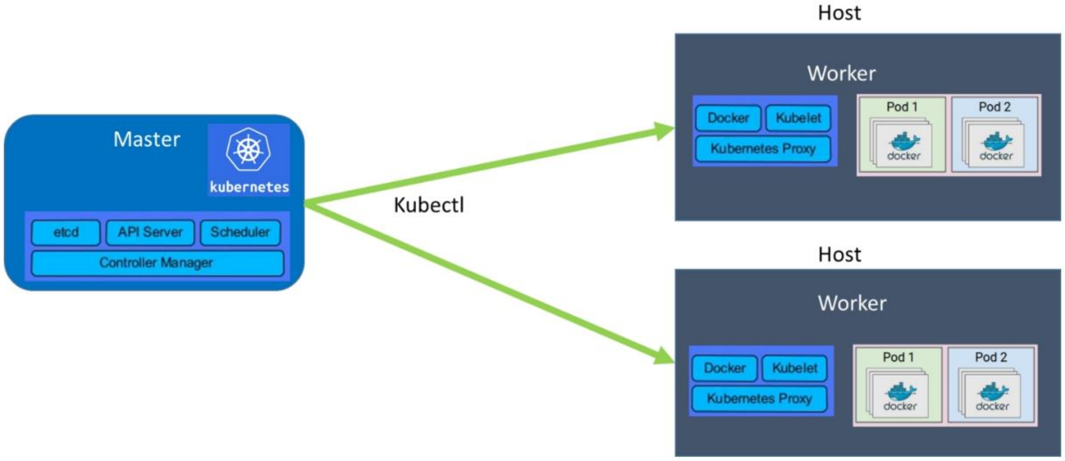

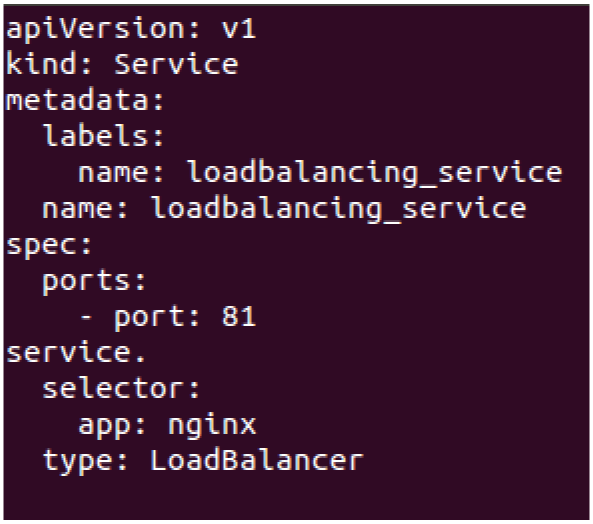

```
apiVersion: v1
kind: Service
metadata:
  labels:
    name: loadbalancing_service
  name: loadbalancing_service
spec:
  ports:
    - port: 81
service.
  selector:
    app: nginx
  type: LoadBalancer
```

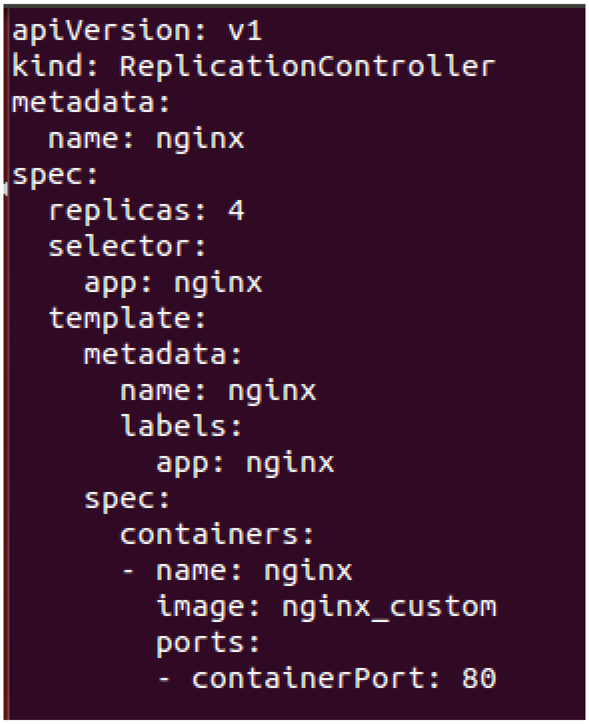

```
apiVersion: v1
kind: ReplicationController
metadata:
  name: nginx
spec:
  replicas: 4
  selector:
    app: nginx
  template:
    metadata:
      name: nginx
      labels:
        app: nginx
    spec:
      containers:
      - name: nginx
        image: nginx_custom
        ports:
        - containerPort: 80
```

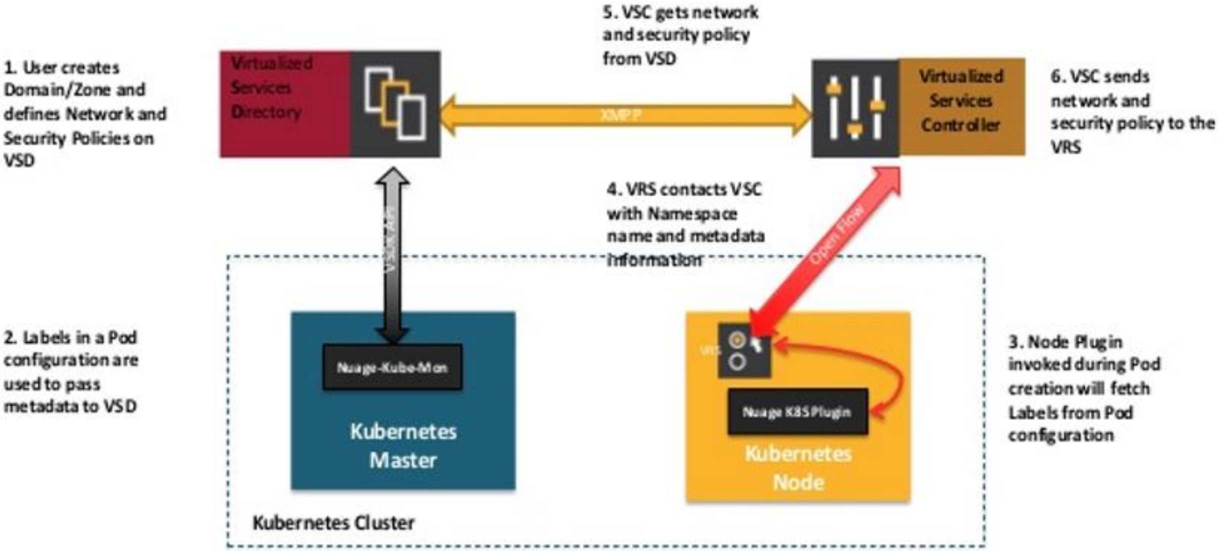

1. User creates Domain/Zone and defines Network and Security Policies on VSD

5. VSC gets network and security policy from VSD

6. VSC sends network and security policy to the VRS

4. VRS contacts VSC with Namespace name and metadata information

2. Labels in a Pod configuration are used to pass metadata to VSD

3. Node Plugin invoked during Pod creation will fetch Labels from Pod configuration

Chapter 11: Securing The Network

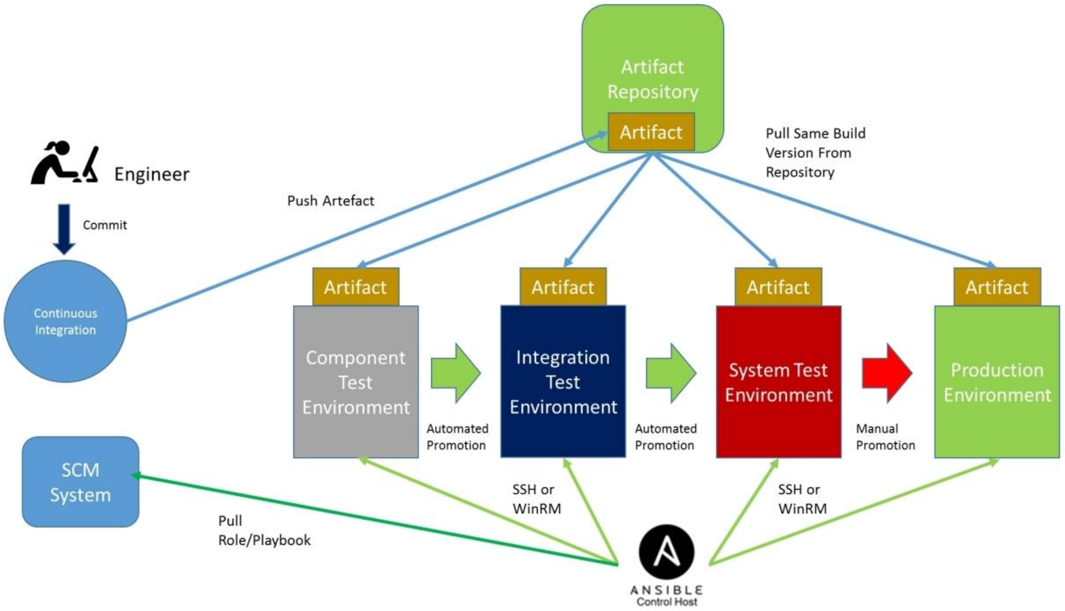

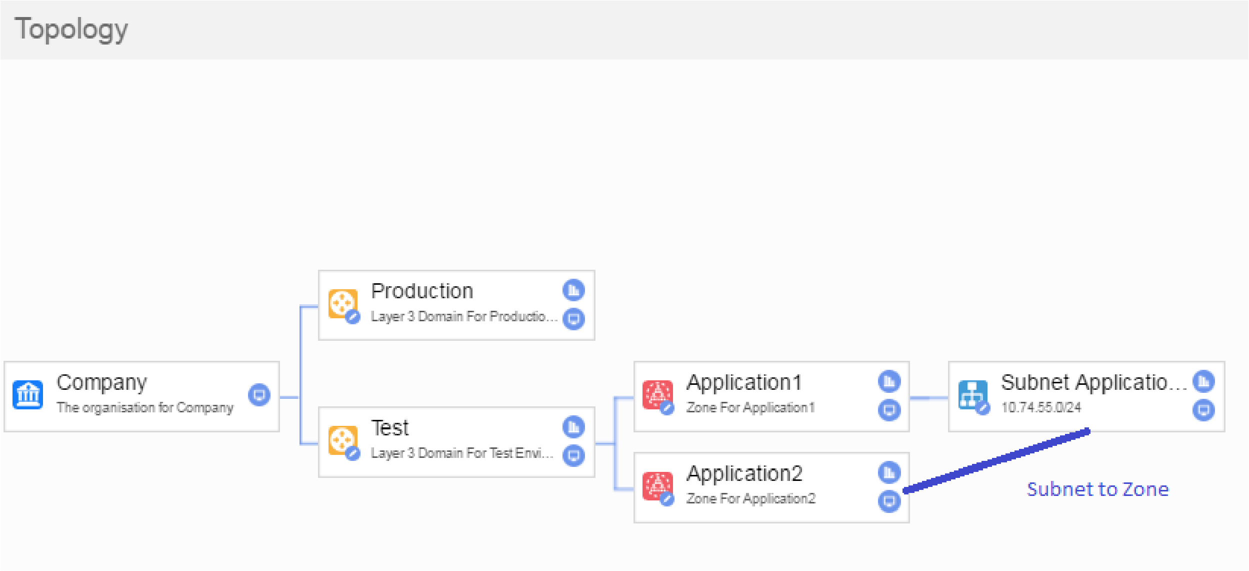

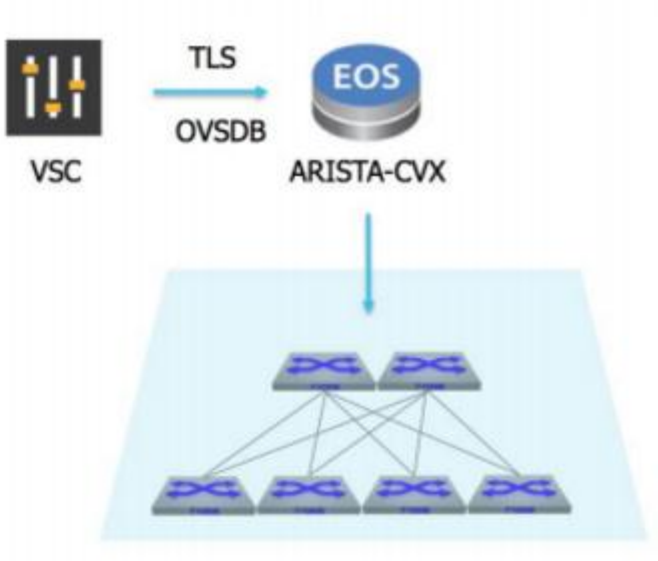

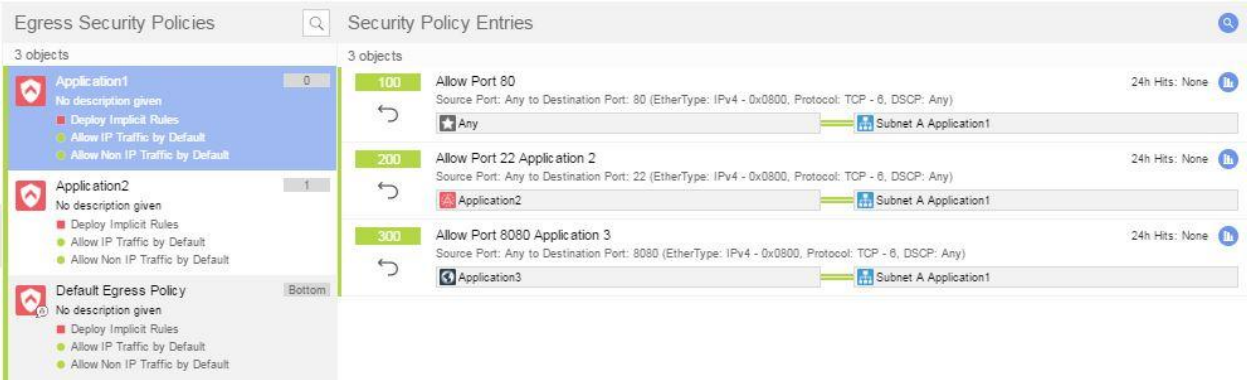

Egress Security Policies

3 objects

Application1 0
No description given
- Deploy Implicit Rules
- Allow IP Traffic by Default
- Allow Non IP Traffic by Default

Application2 1
No description given
- Deploy Implicit Rules
- Allow IP Traffic by Default
- Allow Non IP Traffic by Default

Default Egress Policy Bottom
No description given
- Deploy Implicit Rules
- Allow IP Traffic by Default
- Allow Non IP Traffic by Default

Security Policy Entries

3 objects

100 Allow Port 80 — 24h Hits: None
Source Port: Any to Destination Port: 80 (EtherType: IPv4 - 0x0800, Protocol: TCP - 6, DSCP: Any)
Any — Subnet A Application1

200 Allow Port 22 Application 2 — 24h Hits: None
Source Port: Any to Destination Port: 22 (EtherType: IPv4 - 0x0800, Protocol: TCP - 6, DSCP: Any)
Application2 — Subnet A Application1

300 Allow Port 8080 Application 3 — 24h Hits: None
Source Port: Any to Destination Port: 8080 (EtherType: IPv4 - 0x0800, Protocol: TCP - 6, DSCP: Any)
Application3 — Subnet A Application1

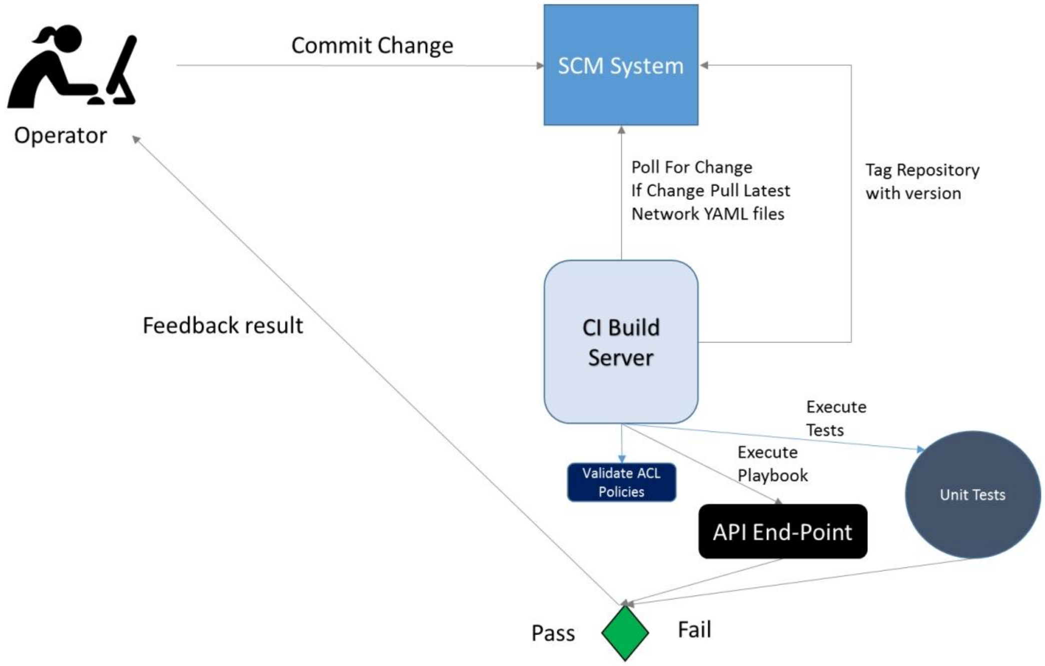

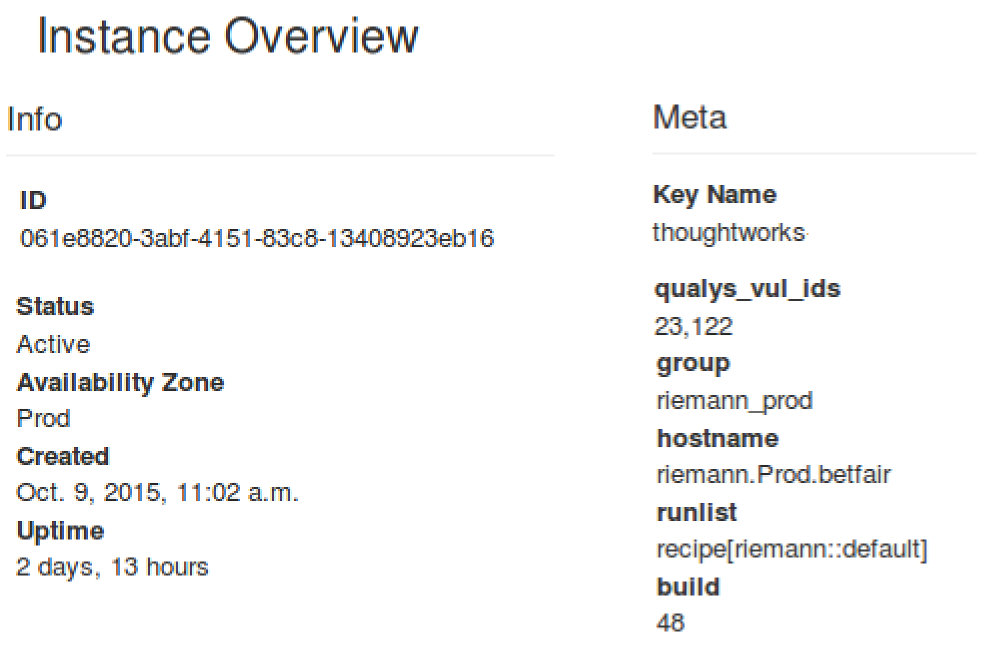

Instance Overview

Info

ID
061e8820-3abf-4151-83c8-13408923eb16

Status
Active

Availability Zone
Prod

Created
Oct. 9, 2015, 11:02 a.m.

Uptime
2 days, 13 hours

Meta

Key Name
thoughtworks

qualys_vul_ids
23,122

group
riemann_prod

hostname
riemann.Prod.betfair

runlist
recipe[riemann::default]

build
48

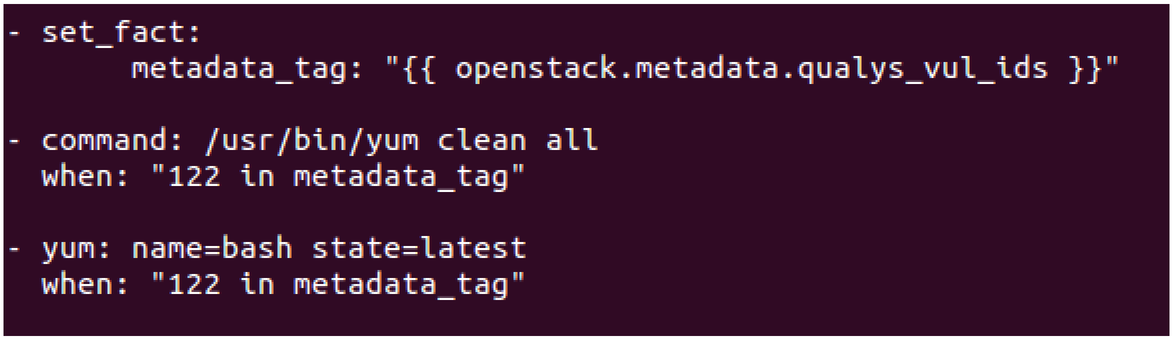

```
- set_fact:
      metadata_tag: "{{ openstack.metadata.qualys_vul_ids }}"

- command: /usr/bin/yum clean all
  when: "122 in metadata_tag"

- yum: name=bash state=latest
  when: "122 in metadata_tag"
```