ICN R5 Test Document

Introduction

ICN Pod Topology

Jenkins Information

Akraino community has a public Jenkins cluster. ICN leverages the Akraino public Jenkins to run CI jobs. While the CD jobs run in our private Jenkins cluster.

We have the following Jenkins slave nodes joined Akraino Jenkins. ICN CI jobs are supposed to be scheduled to our slave nodes by label icn-dev.

Slave Information			**Server Information**
Slave Name	Labels	Slave Root	Server Info
prd-ubuntu-dev-44c-64g	icn-dev	/home/jenkins/akraino/slave_root	pod14-node1

To add more Jenkins slave nodes, please follow the Akriano Jenkins guide

To setup private Jenkins, please refer to the README.md under icn/ci/

The private Jenkins cluster is setup on pod14-node2. We can visit the Jenkins with the node IP address: http://10.10.140.22:8080/

Currently we support only AIO private Jenkins.

Akraino Test Group Information

not applicable

Overall Test Architecture

Test Architecture

We support the following jobs

CI job

- Triggered by gerrit patch creation/update.
- The job runs verify.sh under ICN project. The verify.sh currently has integrated the golang test and bashate test.
- Post +1/-1 for gerrit patch if the build succeeds/fails
- Upload the job log to Nexus server in post-build actions

CD job for test

- Triggered daily automatically (We can also trigger it manually)
- Run a make command, which creates VM(s) and deploys ICN components on the VM(s)
- Upload the job log to Nexus server in post-build actions

CI job detail

Update the verify.sh can update the CI job content.

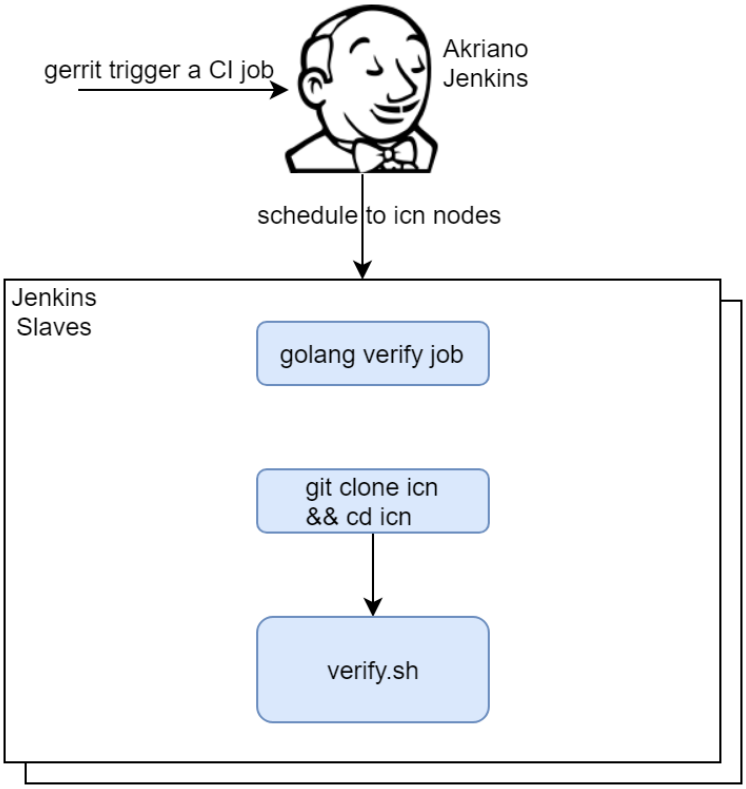

CD job detail

We have the following steps for CD job:

1. On our private Jenkins node, we provision a VM by vagrant. A Vagrantfile which defines the VMs properties is needed. We can define many VM properties in the Vagrantfile:
    - VM hostname
    - VM memory 40G, CPU 32, disk 400GB
2. Login to the VM and run 'make verifier' which installs the components in the VM
3. We destroy the VM as the last step of the job

Test Bed

Pod Topology

ICN Master Bare Metal Deployment Verifier

CD Topology - ICN Master Bare Metal Deployment Verifier

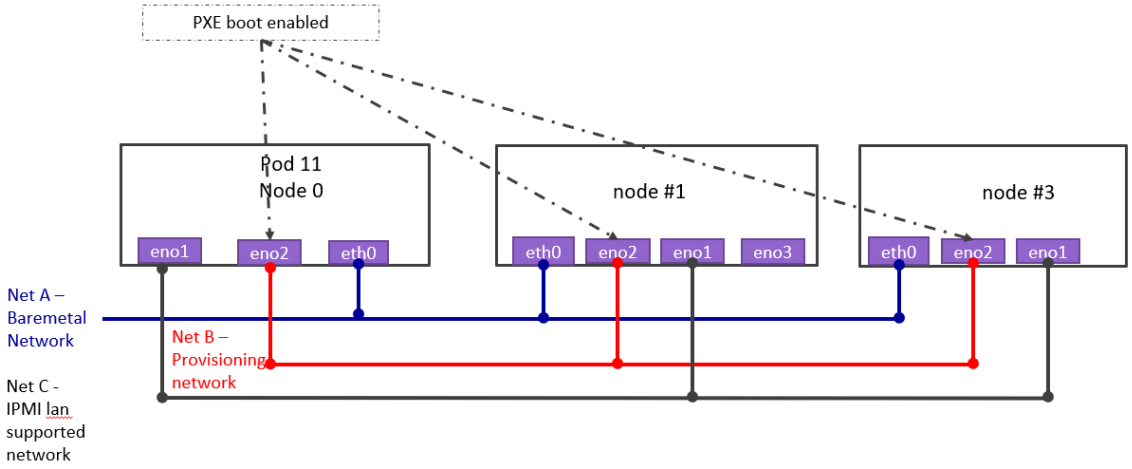

Net A – Baremetal Network, lab networking for SSH, It is used as the control plane for Kubernetes, used by OVN and Flannel for the overlay networking with Internet access
Net B(Internal networking) – Provisioning Networking used by Ironic to do inspection(Will be having DHCP server)
Net C(Internal networking) – IPMI lan to do IPMI protocols for the OS provisioning

ICN Master Virtual Deployment Verifier

CD Topology - ICN Master Virtual Deployment Verifier

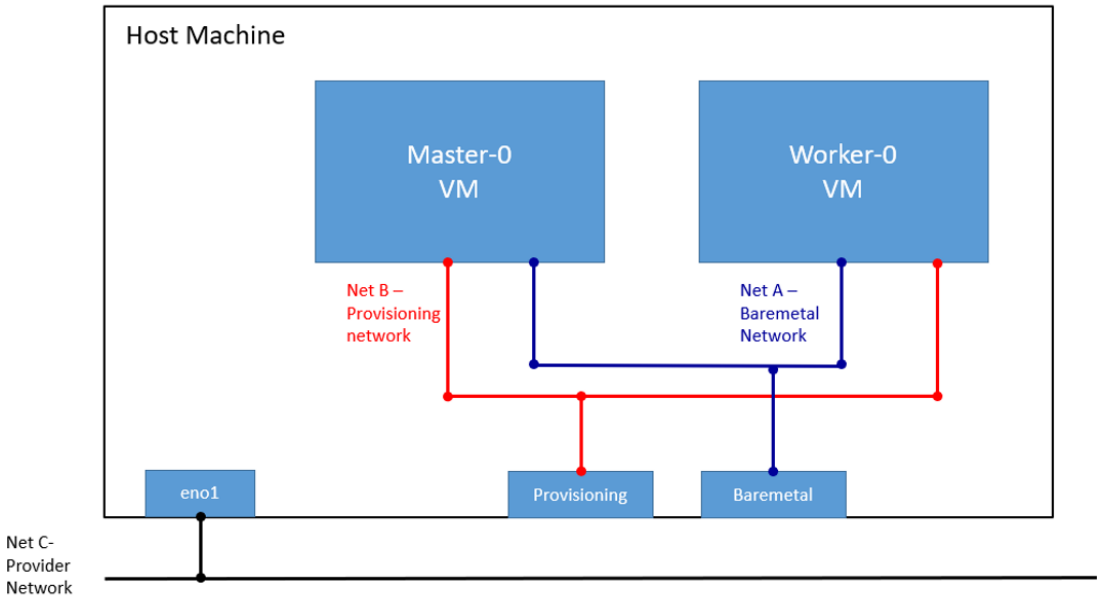

Bare metal deployment

Hostname	CPU Model	Memory	BMC Firmware	Storage	1GbE: NIC#, VLAN, (Connected Extreme 480 switch)	10GbE: NIC# VLAN, Network (Connected with IZ1 switch)	40GbE: NIC#

Jump	Intel 2xE5-2699	64GB	1.46.9995	3TB (Sata) 180 (SSD)	IF0: VLAN 110 (DMZ) IF1: VLAN 111 (Admin)	IF2: VLAN 112 (Private) **VLAN 114** (Management) IF3: VLAN 113 (Storage) **VLAN 1115** (Public)	
node1	Intel 2xE5-2699	64GB	1.46.9995	3TB (Sata) 180 (SSD)	IF0: VLAN 110 (DMZ) IF1: VLAN 111 (Admin)	IF2: VLAN 112 (Private) **VLAN 114** (Management) IF3: VLAN 113 (Storage) **VLAN 1115** (Public)	
node2	Intel 2xE5-2699	64GB	1.46.9995	3TB (Sata) 180 (SSD)	IF0: VLAN 110 (DMZ) IF1: VLAN 111 (Admin)	IF2: VLAN 112 (Private) **VLAN 114** (Management) IF3: VLAN 113 (Storage) **VLAN 1115** (Public)	IF4: SRIOV

Virtual deployment

Hostname	CPU Model	Memory	Storage	1GbE: NIC#, VLAN, (Connected extreme 480 switch)	10GbE: NIC# VLAN, Network (Connected with IZ1 switch)
node1	Intel 2xE5-2699	64GB	3TB (Sata) 180 (SSD)	IF0: VLAN 110 (DMZ) IF1: VLAN 111 (Admin)	IF2: VLAN 112 (Private) **VLAN 114** (Management) IF3: VLAN 113 (Storage) **VLAN 1115** (Public)

Test Framework

All components are tested with end-to-end testing

Test description

Testing

CI Testing:

bashate:

> The bashate test is to check the shell scripts code style. i.e. trailing white space. We find all files with suffix '.sh' and run bashate against the files. '. /cmd/bpa-operator/vendor/' directory is excluded.

golang testing:

BPA Operator:

- The BPA operator has unit tests using the golang framework. The unit tests check the following:
  - Job is created with the right job name for KUD installation.
  - The job metadata has the right cluster name
  - Expected error is produced when a host with the specified MAC address is not found
  - Expected error is produced when no DHCP lease is found for the specified host

BPA Rest Agent:

- Currently, automated unit tests are implemented using the golang testing framework.

CD Verifier (end-to-end testing):

All the test case are tested as follows:

Metal3:

Metal3 verifier will check all the servers are provisioned, Metal3 verifier check the status of the bare metal servers for every 60 second for the provisioning status.

BPA Operator:

Bare Metal Host Provisioning

- The bpa_bmh_verifier.sh script gets the MAC addresses and IP addresses of the 2 VMs provisioned by Metal3, then creates a fake DHCP lease file using the IP address and MAC address information. It also creates a provisioning CR using the MAC address information
- The script the creates an ssh secret key using the ssh keys of the test host, applies the the provisioning CR

- The script busy loops till the KUD installation job completes or fails. If it completes successfully, it does a curl command using the authentication info of the new cluster to confirm if it was successful or not. On completing all the steps, it does a tear down step where it deletes everything it created.

BPA Rest Agent

- Test script, e2e_test.sh, creates dummy image file, creates test JSON file, checks BPA rest agent status, issues POST, GET, and PATCH requests sequentially.
- Next, e2e_test.sh checks uploaded MinIO image object size, and calls DELETE.
- If the script fails at any point then verification was unsuccessful.

Kubernetes Deployment (KUD)

KUD has test cases to verify if the addons are running correctly. All the test cases can be found in tests directory in the multicloud-k8s project. For each of these, we bring up the deployment that is specific to the addon, perform addon specific actions on the pod related to the deployment

Multus:

- Multus CNI is a container network interface (CNI) plugin for Kubernetes that enables attaching multiple network interfaces to pods. This is accomplished by Multus acting as a "meta-plugin", a CNI plugin that can call multiple other CNI plugins.
- A 'NetworkAttachmentDefinition' is used to set up the network attachment, i.e. secondary interface for the pod.
- A pod is created with requesting specific network annotations with bridge CNI to create multiple interfaces. When the pod is up and running, we can attach to it to check the network interfaces on it by running *ip a* command

Nodus:

- Nodus provide Provider networks using VLAN networking and Service Function Chaining.
- After the pod is up and running we will be able to attach to the pod and check for multiple interfaces created inside the container.
- Nodus networking is setup and created

Node Feature Discovery

- Node Feature Discovery for Kubernetes detects hardware features available on each node in a Kubernetes cluster and advertises those features using node labels.
- Create a pod with specific label information in the case the pods are scheduled only on nodes whose major kernel version is 3 and above. Since the NFD master and worker daemonset is already running, the master has all the label information about the nodes which is collected by the worker.
- If the OS version matches, the Pod will be scheduled and up. Otherwise, the Pod will be in a pending state in case there are no nodes with matching labels that are requested by the Pod

SRIOV

- The SRIOV network device plugin is Kubernetes device plugin for discovering and advertising SRIOV network virtual functions (VFs) in a Kubernetes host.
- We first determine which hosts are SRIOV capable and install the drivers on them and run the DaemonSet and register Network Attachment Definition
- On an SRIOV capable hosts, we can get the resources for the node before we run the pod. When we run the test case, there is a request for a VF from the pod, therefore the number of resources for the node is increased.

QAT

- KUD identify if there are QAT devices in the host and decide whether to deploy QAT device plugin into Kubernetes cluster.
- The QAT device plugin discovers and advertises QAT virtual functions (VFs) to Kubernetes cluster.
- KUD assign 1 QAT VF to the Kernel workloads. After the assignment finished, the allocated resources in node description will increase.

CMK

- CPU Manager for Kubernetes provides CPU pinning for K8s workloads. In KUD, there are two test cases for the exclusive and shared CPU pools testing.

SDEWAN

- Use KUD to setup 3 clusters (sdewan-hub, edge-a, edge-b)
- Run the SDEWAN CRD Controller in each clusters.
- Create SDEWAN CNF instance and dummy pod (using httpbin instead) in edge-a, SDEWAN CNF instance and httpbin pod in edge-b
- Create IPSec CR to configure sdewan-hub as responder to provide virtual IP addresses to any authenticated party requesting for IP addresses through SDEWAN CRD Controller.
- Create IPSec CR to configure edge-a and edge-b IPSec configuration to get the IP addresses through SDEWAN CRD Controller.
- Establish edge-a tunnel to sdewan-hub, edge-b tunnel to sdewan-hub, and hub XFRM policies will automatically route traffic between edge-a and edge-b
- Create SNAT CR to establish SNAT rule in edge-a and DNAT CR to establish DNAT rule in edge-b which will enable TCP connection from edge-a to edge-b's httpbin service.
- Verify curl command is successful from edge-a dummy pod (using httpbin instead) to edge-b's httpbin service. The function of the curl command is to return back the ip address of the requester.

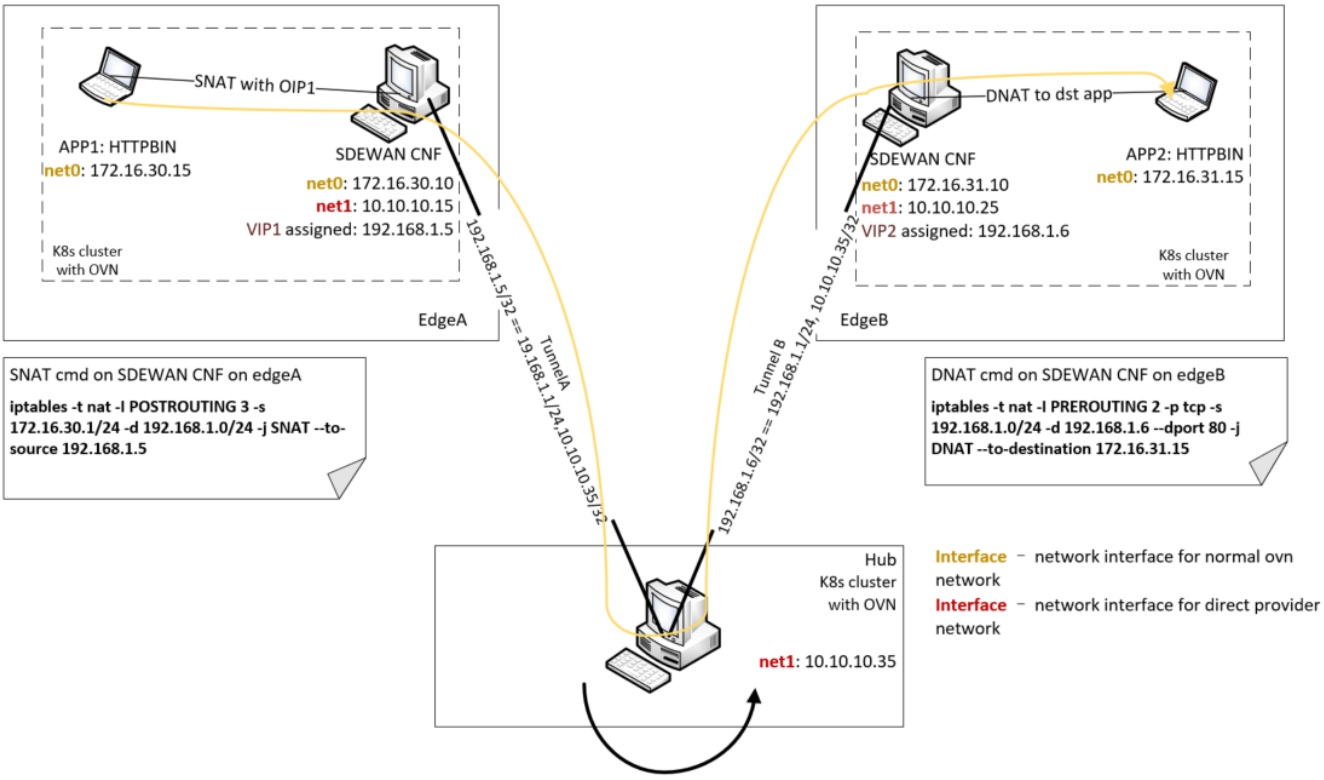

Connection test: Returning back the ip address where it is calling from
root@simple-http-service-84b4b4ccc9-6xtxt:/# curl -X GET "http://192.168.1.6/ip" -H "accept: application/json"
{
"origin": "192.168.1.5"
}

EMCO:

- EMCO sanity testing check the health connectivity to EMCO micro services, once it is installed

BluVal Testing

Status as of July 7th 2021:

Layer	Result	Comments	Nexus
os/lynis	PASS with exceptions	Exceptions: ▪ USB-2000 ▪ SSH-7408: Checking MaxSessions, Checking Port ▪ KRNL-6000: net.ipv4.conf.all.forwarding	Logs
os/vuls	PASS with exceptions	Exceptions: • CVE-2016-1585 • CVE-2017-18342 • CVE-2017-8283 • CVE-2018-20839 • CVE-2019-17041 • CVE-2019-17042 • CVE-2019-19814	Logs
k8s/conformance	PASS with exceptions	Exceptions: • Sonobuoy v0.16.1 does not support Kubernetes v1.18.9	Logs
k8s/kube-hunter	PASS	With aquasec/kube-hunter:edge image	Logs

Release 5 Blueprint Scanning Status

Akraino CVE Vulnerability Exception Request

Akraino BluVal Exception Request

CD logs:

ICN Master Bare Metal Deployment Verifier

ICN Master Virtual Deployment Verifier

ICN SDEWAN Master End2End Testing

Test Dashboards

All the testing results are in logs

Additional Testing

not applicable

Bottlenecks/Errata

not applicable