

Default System Configurations

LACP

In Cisco SD-WAN Cloud OnRamp for Colocation solution, the Link Aggregation Control Protocol (LACP) is enabled for the management port channel. The management port channel is created by default using the Ethernet links (eth0-1 and eth0-2). To ensure the port channel configuration on the management switch side is reachable, run the **support ovs appctl bond-show mgmt-bond** command and ensure OOB switch ports that are connected to the switch has the following port-channel configuration.

```
!
interface Port-channel1
 switchport mode access
!
interface GigabitEthernet1/0/6
 switchport mode access
 channel-group 1 mode passive
!
interface GigabitEthernet1/0/7
 switchport mode access
 channel-group 1 mode passive
!
```

DHCP

In Cisco SD-WAN Cloud OnRamp for Colocation solution, DHCP is enabled by default on the management port channel. Once the DHCP server is up, the host gets the DHCP IP address on management internal port.

Sticky DHCP

Note Sticky DHCP configurations are optional.

Configure the DHCP servers to get sticky DHCP IP address. The DHCP client identifier is the serial number on the CSP device.

The DHCP server configuration on a Linux server is:

```
host
{
option dhcp-client-identifier "WZP22060AUR";
fixed-address 10.20.0.2;
```

```
option routers 10.20.0.1;
option domain-name-servers 198.51.100.9;
option domain-name "cisco.com";
option subnet-mask 255.255.0.0;
}
```

The DHCP server configuration on IOS is:

```
ip dhcp pool P_112
host 10.0.0.2 255.255.0.0
client-identifier 4643.4832.3133.3256.3131.48
default-router 10.0.0.1
dns-server 10.0.0.1
```

Here 10.0.0.2 is the sticky DHCP IP address. Use **debug ip dhcp server packet** command to find out the client identifier.

Static IPv4

To troubleshoot issues with DHCP configurations, configure static IPv4 on the management port channel:

```
configure shared
vm_lifecycle networks network int-mgmt-net subnet int-mgmt-net-subnet address 105.20.0.0
gateway 105.20.0.1 netmask 255.255.255.0 dhcp false
system settings domain cisco.com
system:system settings dns-server 209.165.201.20
system:system settings ip-receive-acl 0.0.0.0/0
action accept
priority 100
service scpd
commit
```

Since vManage is the controller in this solution, **Configure shared** writes to candidate database (CDB) which will keep the device config in sync with vManage.

Note **Configure shared** is only applicable to static ip configurations. Any other configurations done manually using either confd cli or netconf or rest api, will be removed by vManage as NFVIS is a vManaged device in this solution.

In NFVIS NetworkHub image, networks are automated and user should not create, delete or modify networks. You can reset the host server to default configurations using the **factory reset all** command.

- SRIOV Support, on page 2
- OVS-DPDK Support, on page 3

SRIOV Support

SR-IOV is statically enabled on NFVIS Cisco SD-WAN Cloud OnRamp for Colocation image with a CSP 5444 Product Identifier (PID).

- SRIOV is enabled by default on ethernet ports eth1-1, eth1-2, eth4-1, eth4-2 as Niantec NIM cards are placed in slots 1 and 4.
- SR-IOV is enabled only on Niantic NICs and onboard Niantics does not support SR-IOV.

- Thirty two virtual functions are created on each PNIC. . If the NIC is connected to 1G, two virtual functions are created.
- Virtual Ethernet Port Aggregator (VEPA) mode is enabled.
- The naming convention is: <interface name>-SRIOV-1,<interface name>-SRIOV-2 ,<interface name>-SRIOV-3,<interface name>-SRIOV-4.
- Fortville NICs are used to create port channels for OVS data traffic and HA sync between the VM's.

OVS-DPDK Support

Starting from NFVIS 3.12 release, OVS-DPDK support is enabled on NFVIS. The data and HA bridges are ovs-dpdk bridges by default. The bridges are associated with bonds created over pnics on the Fortville card. The pnics are bound to DPDK compatible drivers. OVS-DPDK support provides higher performance than the standard kernel OVS datapath.

NFVIS System has two CPUs from the same socket. For each socket one CPU is reserved for DPDK. The kernel is allocated 8GB of memory. 4G memory is allocated per socket for the DPDK poll mode driver. The rest of the memory is converted to 2MB huge pages and allocated for VM deployment.

To check the OVS-DPDK status, use **Show system:system settings dpdk-status** command.