VMware Cloud Director Availability Multi-NIC Setup

Steps for configuring each of the appliances

Authors:
Atanas Stankov, Senior Solutions Architect, VMware
Nikolay Patrikov, Senior TPM, VMware

vmware

Table of contents

Overview

Placing the VMware Cloud Director Availability appliances in different networks is something common for Service Providers. This leads to the need to use multiple network interfaces to support that scenario and guarantee the product will continue operating as per design.

There are a few main reasons that lead to this design:

- Bypass a router
- Port forwarding is not possible
- Inability to route the replication traffic
- Separation of the incoming replications to different isolated networks

Design

You need to take two considerations before the implementation of VMware Cloud Director Availability appliances that will have multiple NICs:

- Which interface will be used for the communication with the rest of the VMware Cloud Director Availability appliances?
- How will the routing be organized - to which interface the default gateway will be configured and what static routes will be required?

The **General recommendation (GR01)** is to deploy the appliances with the interfaces that will be used for communicating with each other. When following this recommendation, the VMware Cloud Director Availability services will discover and set to use the first NIC (ens160) and its first IP address. The only additional change that might be required is to move the default gateway to a different NIC and configure one or more static routes on ens160.

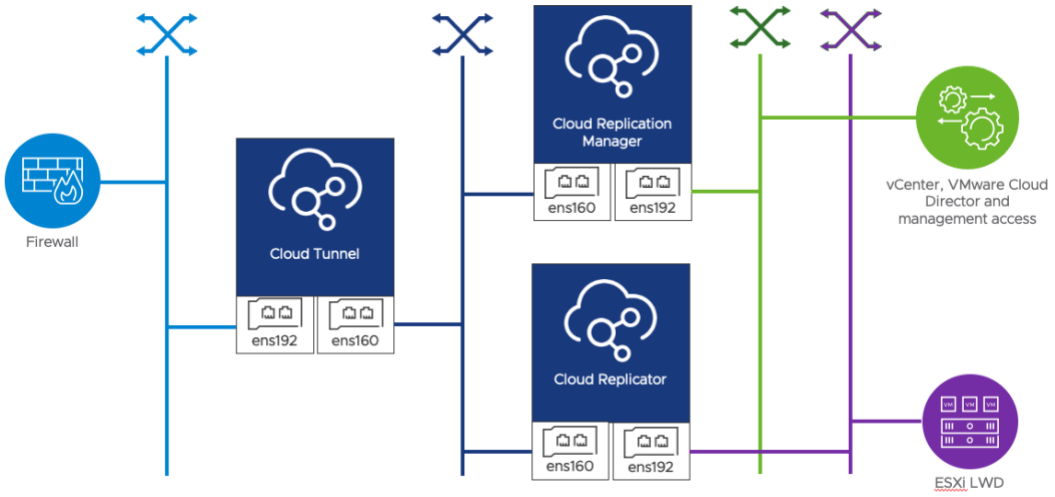

Figure 1 - Example network diagram

Moving the Default Gateway

To perform the necessary steps, you have to connect to the appliance. The recommended way is to use DCUI because the configuration changes might lead to losing network connectivity.

1. First, you need to unconfigure ens160. Example command:

   /opt/vmware/h4/bin/net.py unconfigure-nic ens160

2. Then configure only the IP address of ens160 without setting a default gateway on it. Example command:

   /opt/vmware/h4/bin/net.py configure-nic --static -a 192.168.10.10/24 ens160

3. Configure the static route(s) for ens160. Example command:

   /opt/vmware/h4/bin/net.py add-route ens160 destination_ subnet gateway metric

   Note: You can use this command multiple times for setting more than one static route.

4. (Optional) Configure the name servers. Example command:

   `/opt/vmware/h4/bin/net.py configure-dns --servers ns01_ip ns02_ip --search-domain domain1 domain2`

Once you finish the configuration and the appliance is accessible over the network, you can connect through SSH or use the UI.

Reconfigure the VMware Cloud Director Availability Appliances

For cases where **GR01** is not followed, some actions need to be taken so the VMware Cloud Director Availability appliances can use the IP address set to an interface different from ens160.

Cloud Tunnel

It is essential to know that only one of the Cloud Tunnel interfaces can be used for communication with the rest of the VMware Cloud Director Availability appliances in the local site. For pairing purposes all the interfaces of the tunnel appliance can be used.

To reconfigure the Cloud Tunnel, you need to:

1. Set IP addresses to all interfaces and configure the static routes.
2. Log in as **root** through SSH to the Tunnel appliance.
3. Authenticate as **root** to the Tunnel service via CLI. Example command:

   `h4 -k tunloginroot 'r00t_Password'`

4. Check the current service configuration. Example command:

   `h4 -k tunendpoints`

   Expected output:


```
{
  "configured": {
    "mgmtAddress": "192.168.1.2",
    "mgmtPort": 8047,
    "mgmtPublicAddress": "192.168.1.2",
    "mgmtPublicPort": 8047,
    "tunnelAddress": "192.168.1.2",
    "tunnelPort": 8048,
    "tunnelPublicAddress": "vcav01.ber.cloudprovider.pub",
    "tunnelPublicPort": 443
  },
  "effective": {
    "mgmtAddress": "192.168.1.2",
    "mgmtPort": 8047,
    "mgmtPublicAddress": "192.168.1.2",
    "mgmtPublicPort": 8047,
    "tunnelAddress": "192.168.1.2",
    "tunnelPort": 8048,
    "tunnelPublicAddress": "vcav01.ber.cloudprovider.pub",
    "tunnelPublicPort": 443
  }
}
```

The important parameter is tunnelAddress. The value of this parameter will be used to configure the other local VMware Cloud Director Availability appliances – Cloud Replication Manager and Cloud Replicator(s) - when they are prepared to communicate with the tunnel. This parameter can be set to a specific IP address or null. Setting it to null will lead to the service discovering the first IP address in the system and use it. In this example, the **ens160** IP address is 192.168.1.2, and the **ens192** IP address is 172.18.24.4.

5. Next, we need to configure the Tunnel service to use the IP address of **ens192**. Example command:

   h4 -k tunsetendpoints "192.168.1.2" "8047" "192.168.1.2" "8047" "172.18.24.4" "8048" "vcav01.ber.cloudprovider.pub" "443"

   Expected output:


```
{
  "configured": {
    "mgmtAddress": "192.168.1.2",
    "mgmtPort": 8047,
    "mgmtPublicAddress": "192.168.1.2",
    "mgmtPublicPort": 8047,
    "tunnelAddress": "172.18.24.4",
    "tunnelPort": 8048,
    "tunnelPublicAddress": "vcav01.ber.cloudprovider.pub",
    "tunnelPublicPort": 443
  },
  "effective": {
    "mgmtAddress": "192.168.1.2",
    "mgmtPort": 8047,
    "mgmtPublicAddress": "192.168.1.2",
    "mgmtPublicPort": 8047,
    "tunnelAddress": "172.18.24.4",
    "tunnelPort": 8048,
    "tunnelPublicAddress": "vcav01.ber.cloudprovider.pub",
    "tunnelPublicPort": 443
  }
}
```


   Note: Make sure the parameters are provided in the correct order.

With this, the Cloud Tunnel is properly configured to use a non-default IP address.

Once the **tunnelAddress** has its new value, please navigate to the VMware Cloud Director Availability Portal (*https://manager_appliance_IP_address/ui/admin*) and re-enable the tunneling to propagate the change to all local components – Cloud Replication Manager and Cloud Replicator(s).

Cloud Replication Manager

The Cloud Replication Manager runs two services (Cloud service and Manager service) that require reconfiguration as both need to use the same interface to communicate with the rest of the appliances. For communication with the VMware Cloud Director cells and vCenters, the services can use any interface.

Cloud service

1. Set IP addresses to all interfaces and configure the static routes.
2. Log in as **root** through SSH to the Cloud Replication Manager appliance.
3. Authenticate as **root** to the Cloud service via CLI. Example command:

   `c4 loginroot 'r00t_Password'`
4. Check the current service configuration. Example command:

   `c4 endpoints`

   Expected output:


```
{
  "configured": {
    "mgmtAddress": null,
    "mgmtPort": 8046,
    "mgmtPublicAddress": "tn-e2eecba4-8381-4799-a39a-0ec4eca105bb.tnexus.io",
    "mgmtPublicPort": 8048,
    "apiAddress": null,
    "apiPort": 8443,
    "apiPublicAddress": "vcav01.ber.cloudprovider.pub",
    "apiPublicPort": 443
  },
  "effective": {
    "mgmtAddress": "192.168.2.81",
    "mgmtPort": 8046,
    "mgmtPublicAddress": "tn-e2eecba4-8381-4799-a39a-0ec4eca105bb.tnexus.io",
    "mgmtPublicPort": 8048,
    "apiAddress": "vcavm1.ber.cloudprovider.local",
    "apiPort": 8443,
    "apiPublicAddress": "vcav01.ber.cloudprovider.pub",
    "apiPublicPort": 443
  }
}
```


   Note: Because the Cloud Replication Manager appliance is already prepared for tunneling, the **mgmtPublicAddress** parameter has this value **tn-<<uuid>>.tnexus.io**. It should NOT be changed in the next step but should be used as it is.

   The parameter that has to be changed is **mgmtAddress**. In the Configured section it is set to **null** which means the service will try to discover an interface with a successfully configured IP address and bind itself to this address.

   In this example, the **ens160** IP address is 192.168.2.81, and the service currently uses this interface. The IP address of the **ens192** interface is 172.17.31.81.
5. We will configure **mgmtAddress** to use ens192 with IP address 172.17.31.81. Example command:

   `c4 -k setendpoints "172.17.31.81" "8046" "tn-e2eecba4-8381-4799-a39a-0ec4eca105bb.tnexus.io" "8048" "vcavm1.ber.cloudprovider.local" "8443" "vcav01.ber.cloudprovider.pub" "443"`

Expected output:


```
{
  "configured": {
    "mgmtAddress": "172.17.31.81",
    "mgmtPort": 8046,
    "mgmtPublicAddress": "tn-e2eecba4-8381-4799-a39a-0ec4eca105bb.tnexus.io",
    "mgmtPublicPort": 8048,
    "apiAddress": "vcavm1.ber.cloudprovider.local",
    "apiPort": 8443,
    "apiPublicAddress": "vcav01.ber.cloudprovider.pub",
    "apiPublicPort": 443
  },
  "effective": {
    "mgmtAddress": "172.17.31.81",
    "mgmtPort": 8046,
    "mgmtPublicAddress": "tn-e2eecba4-8381-4799-a39a-0ec4eca105bb.tnexus.io",
    "mgmtPublicPort": 8048,
    "apiAddress": "vcavm1.ber.cloudprovider.local",
    "apiPort": 8443,
    "apiPublicAddress": "vcav01.ber.cloudprovider.pub",
    "apiPublicPort": 443
  }
}
```


6. Restart the Cloud service. Example command:

   systemctl restart manager

   Note: This operation won't break any of the running replication tasks but will only reload the Cloud Replicator configuration and ensure all paired on-premises and remote cloud appliances receive the correct information for the replicator configuration in the local site.

   Once you perform these steps, the Cloud service is configured, and you can proceed to the steps for the Manager service.

Manager service

1. Log in as **root** through SSH to the Cloud Replication Manager appliance.
2. Authenticate as **root** to the Manager service via CLI. Example command:

   h4 loginroot 'r00t_Password'
3. Check the current service configuration. Example command:

   h4 endpoints

Expected output:


```
{
  "configured": {
    "mgmtAddress": null,
    "mgmtPort": 8044,
    "mgmtPublicAddress": "tn-853bd005-b905-4da1-bb75-e17996efb5df.tnexus.io",
    "mgmtPublicPort": 8048
  },
  "effective": {
    "mgmtAddress": "192.168.2.81",
    "mgmtPort": 8044,
    "mgmtPublicAddress": "tn-853bd005-b905-4da1-bb75-e17996efb5df.tnexus.io",
    "mgmtPublicPort": 8048
  }
}
```


Note: The **mgmtAddress** has **null** value. This is the parameter that will be changed in the next step.

4. Change the **mgmtAddress**. Example command:

   h4 setendpoints "172.17.31.81" "8044" "tn-853bd005-b905-4da1-bb75-e17996efb5df.tnexus.io" "8048"

   Expected output:


```
{
  "configured": {
    "mgmtAddress": "172.17.31.81",
    "mgmtPort": 8044,
    "mgmtPublicAddress": "tn-853bd005-b905-4da1-bb75-e17996efb5df.tnexus.io",
    "mgmtPublicPort": 8048
  },
  "effective": {
    "mgmtAddress": "172.17.31.81",
    "mgmtPort": 8044,
    "mgmtPublicAddress": "tn-853bd005-b905-4da1-bb75-e17996efb5df.tnexus.io",
    "mgmtPublicPort": 8048
  }
}
```


With this, the configuration of the Replication Manager is completed.

To restore the regular operation of VMware Cloud Director Availability, you will need to:

- Re-register all the Cloud Replicators via the Manager service UI (*https://manager_appliance_IP_address:8441*)
- Enable tunneling via the VMware Cloud Director Availability Portal (*https://manager_appliance_IP_address/ui/admin*)

Cloud Replicator(s)

The idea behind having multiple interfaces for the Cloud Replicator is to optimize replication traffic flow. When the Cloud Replicator has more than one NIC, one of the interfaces should be able to communicate with the rest of the VMware Cloud Director Availability appliances and the other one should be in the same Layer 2 broadcast domain as the ESXi vmkernel interface. The Cloud Replicator can't use one interface to communicate with the Cloud Tunnel and another interface to communicate with the Cloud Replication Manager.

1. Set IP addresses to all interfaces and configure the static routes.
2. Log in as **root** through SSH to the Cloud Replicator appliance.
3. Authenticate as **root** to the Replicator service via CLI. Example command:

   h4 rtrloginroot 'r00t_Password'
4. Check the current service configuration. Example command:

   h4 rtrendpoints

   Expected output:


```
{
  "configured": {
    "mgmtAddress": null,
    "mgmtPort": 8043,
    "mgmtPublicAddress": "tn-f9abd9de-3978-4383-a951-514deaec522f.tnexus.io",
    "mgmtPublicPort": 8048,
    "nfcAddress": null,
    "lwdAddress": null,
    "lwdPort": null,
    "lwdPublicAddress": "lw-f9abd9de-3978-4383-a951-514deaec522f.tnexus.io",
    "lwdPublicPort": 8048
  },
  "effective": {
    "mgmtAddress": "172.17.33.1",
    "mgmtPort": 8043,
    "mgmtPublicAddress": "tn-f9abd9de-3978-4383-a951-514deaec522f.tnexus.io",
    "mgmtPublicPort": 8048,
    "nfcAddress": null,
    "lwdAddress": "172.17.33.1",
    "lwdPort": 44045,
    "lwdPublicAddress": "lw-f9abd9de-3978-4383-a951-514deaec522f.tnexus.io",
    "lwdPublicPort": 8048
  }
}
```


Note: Because the Cloud Replicator appliance is already prepared for tunneling, the **mgmtPublicAddress** parameter has this value **tn-<<uuid>>.tnexus.io**. It should NOT be changed in the next step but should be used as it is.

Two parameters require changing:

- **mgmtAddress** – this is the IP address of the interface that will be used for communicating with the other VMware Cloud Director Availability appliances in the local site
- **lwdAddress** – this is the IP address of the interface that will be used for communicating with the replication vmkernel interface of the ESXi host(s).

5. Change the **mgmtAddress** and **lwdAddress**. Example command:

   h4 rtrsetendpoints "172.17.33.1" "8043" "tn-f9abd9de-3978-4383-a951-514deaec522f.tnexus.io" "8048" "" "192.168.3.1" "lw-f9abd9de-3978-4383-a951-514deaec522f.tnexus.io" "8048"

   Note: The parameters need to be in this exact order:

   h4 rtrsetendpoints **mgmtAddress** mgmtPort mgmtPublicAddress mgmtPublicPort nfcAddress **lwdAddress** lwdPublicAddress lwdPublicPort

   Expected output:


```
{
  "configured": {
    "mgmtAddress": "172.17.33.1",
    "mgmtPort": 8043,
    "mgmtPublicAddress": "tn-f9abd9de-3978-4383-a951-514deaec522f.tnexus.io",
    "mgmtPublicPort": 8048,
    "nfcAddress": null,
    "lwdAddress": "192.168.3.1",
    "lwdPort": null,
    "lwdPublicAddress": "lw-f9abd9de-3978-4383-a951-514deaec522f.tnexus.io",
    "lwdPublicPort": 8048
  },
  "effective": {
    "mgmtAddress": "172.17.33.1",
    "mgmtPort": 8043,
    "mgmtPublicAddress": "tn-f9abd9de-3978-4383-a951-514deaec522f.tnexus.io",
    "mgmtPublicPort": 8048,
    "nfcAddress": null,
    "lwdAddress": "192.168.3.1",
    "lwdPort": 44045,
    "lwdPublicAddress": "lw-f9abd9de-3978-4383-a951-514deaec522f.tnexus.io",
    "lwdPublicPort": 8048
  }
}
```


   Note: You might get the following output:


```
"code": "TrafficIsolationConfigurationFailure",
  "msg": "Traffic isolation configuration could not be applied to hbr server.",
  "args": [],
  "stacktrace":
```


   In such cases, just check if the configuration has been updated correctly using:

   h4 rtrendpoints

With this, the configuration of the Cloud Replicator is completed.

When the **mgmtAddress** parameter is changed, the Cloud Replicator needs to be re-registered in the Cloud Replication Manager via the Manager service UI (*https://manager_appliance_IP_address:8441*).

Upgrade

Upgrading the VMware Cloud Director Availability appliances will not affect the network interface configuration.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 vmware.com