

Deployment Guide for VMware vCloud Director

Virtual Multimedia Resource Function

Installation Instructions

Copyright

© Ericsson AB 2018, 2019. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

Disclaimer

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

Trademark List

All trademarks mentioned herein are the property of their respective owners. These are shown in the document *Trademark Information*.



Contents

1	About This Document	1
2	vMRF Deployment Principles for VMware vCloud Director	2
3	vMRF Deployment Process for VMware vCloud Director	3
4	Prerequisites for vMRF Deployment	6
4.1	Configure vCloud Director Organization	6
4.2	Configure vCloud Director to Assign Unique UUIDs	6
4.3	Download and Extract vMRF Software Delivery Package	7
5	vMRF Deployment Preparations for the Cloud Administrator	8
5.1	Cloud Hardware and Software Preparation and Configuration	8
5.2	Create Network Topology	8
6	vMRF Deployment for the End User	10
6.1	Initial VNF Configuration Data for Deployment	10
7	Configure Shared Storage	11
8	Onboard to the Catalog	13
9	Deploy vApp from the Catalog	14
10	Provide Initial Configuration from VNF Configuration ISO File	17
11	Power on vMRF vApp	18
12	Scaling Out to Full VNF Size	19
12.1	Scale-out Using the Catalog	19
13	Check vMRF Status	20





1 About This Document

This document describes vMRF deployment on a VMware cloud service. VMware service means VMware vCloud Director® including VMware ESXi® and VMware vCenter Server®.

The following user roles are distinguished in this document:

Cloud Administrator

The cloud administrator is the cloud service provider who delivers the cloud service to the end user. The cloud administrator must fulfill certain prerequisites before the end user can start deploying vMRF.

End User

The end user is the vMRF operator and deployment responsible, who is assumed to be a cloud service consumer on a vCloud cloud service. The end user is also referred to as a tenant.

2 vMRF Deployment Principles for VMware vCloud Director

If the hardware and software requirements are met, and after the needed configurations in VMware are done, vMRF is instantiated.

vMRF can contain one or more Virtual Network Functions (VNF).

A single VNF contains multiple Virtual Machines (VMs). See [Figure 1](#) for an example overview of vMRF deployment with two VNFs.

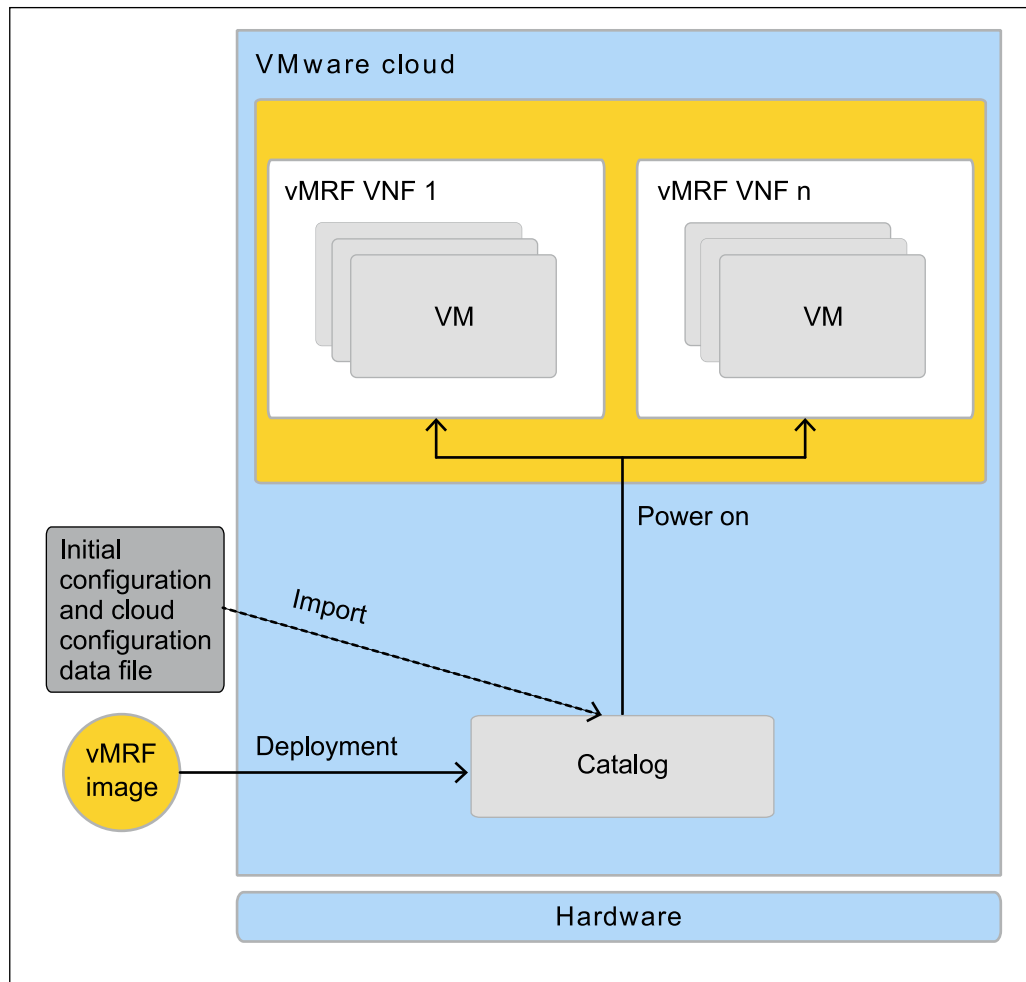


Figure 1 vMRF Deployment



3 vMRF Deployment Process for VMware vCloud Director

The vMRF deployment process consists of preparations and basic configuration of the cloud environment, and the actual instantiation of one or more vMRF VNF instances.

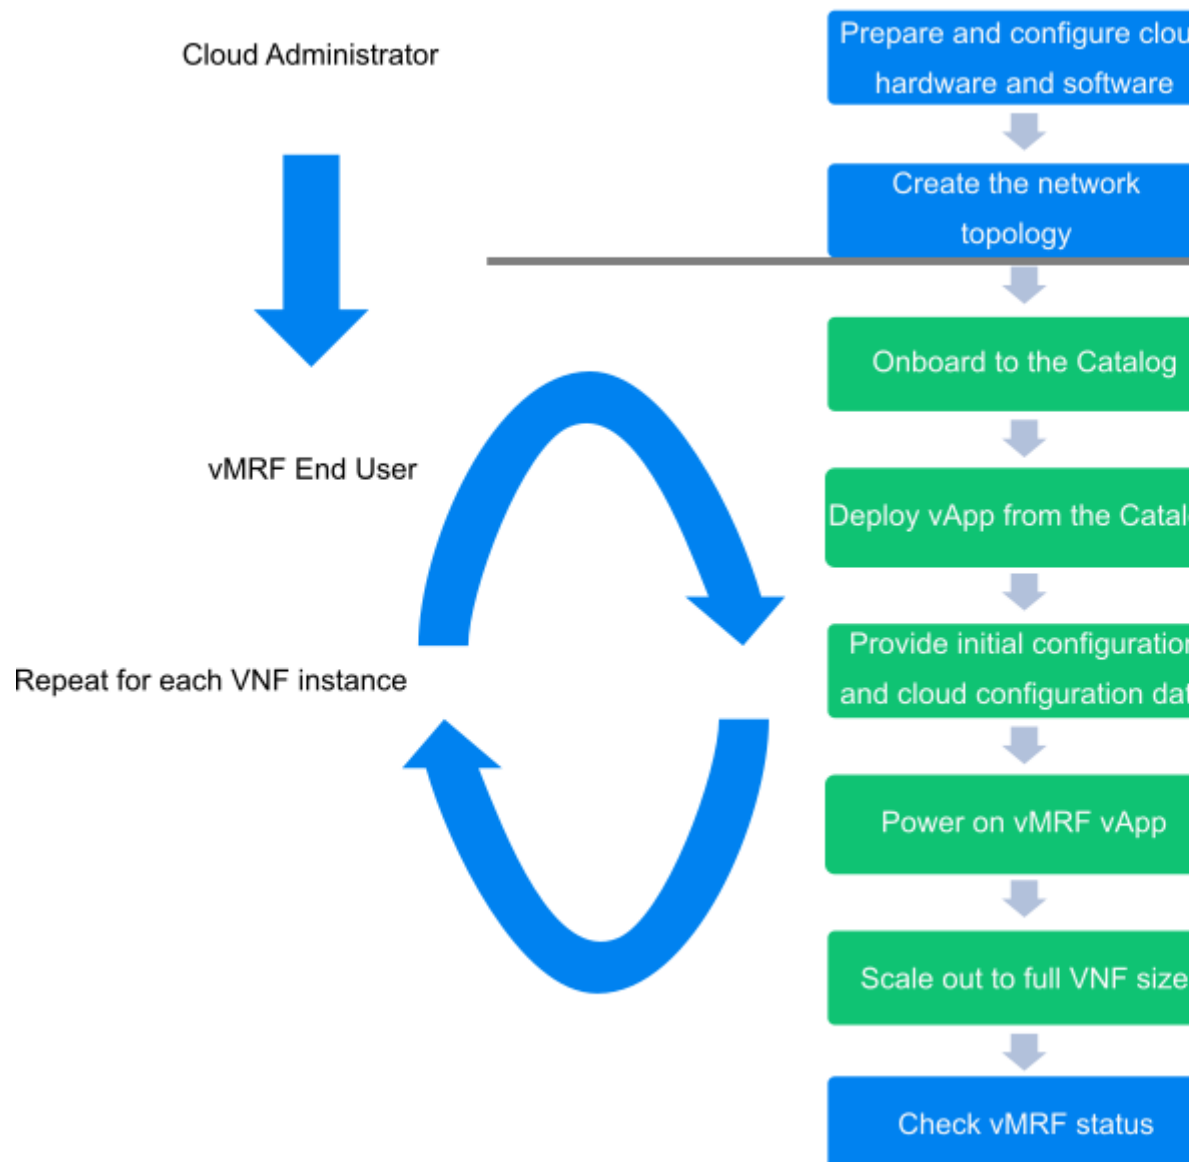


Figure 2 vMRF Deployment Process



1. Prepare the cloud environment to run vMRF

This set of steps is done by the cloud administrator.

a. Prepare and configure cloud hardware and software

This step involves checking that the necessary hardware exists, and making hardware-related configuration in VMware and in the host Operating System so that the requirements listed in [Prerequisites for vMRF Deployment](#) on page 6 are fulfilled.

b. Create the network topology

This step involves ensuring that the required networks to which the VNF connects are in place. It requires the cloud administrator to set up the networks, create distributed port groups in vSphere and external networks in vCD; and organization administration to create organization specific networks and map them to the external networks.

2. Deploy and check vMRF

This set of steps is done by the end user.

a. Download and extract the vMRF software delivery package

The vMRF software delivery package contains the Open Virtualization Format (OVF) template and the related files. The vMRF software delivery package must be extracted to a place where the files can be accessed by the vCloud Director. vCloud Director offers the following options:

- A local directory that is accessible by the vCloud Director client
- The files can be uploaded to an HTTP server accessible from the vCloud Director

b. Onboard the software delivery package to the Catalog.

c. Deploy vApp from the Catalog and Provide Initial Configuration and Cloud Configuration Data

d. Power On vMRF vApp

After powering on the vMRF vApp, if the initial configuration is provided, and the components are operating with normal status, the VNF is capable of running traffic.

It is recommended to check the status of the powered-on vMRF vApp before scaling out to full VNF size.

e. Scale-out to Full VNF Size

f. Check vMRF status



It is recommended to run a status check on the newly deployed vMRF.



4 Prerequisites for vMRF Deployment

Before the end user can deploy and use vMRF, the cloud administrator must ensure that the environment fulfills hardware, software, and network requirements. The main requirements are listed in [vMRF Infrastructure Requirements](#).

4.1 Configure vCloud Director Organization

The vCD organization must be prepared for the use of the following ExtraConfig OVF parameters to avoid degraded performance in high traffic load situations:

- ethernet <X>.ctxPerDev
- sched.mem.pin
- sched.cpu.latencySensitivity

For more information on setting up vCD organization environment, see [vCloud API Programming Guide for Service Providers](#).

Steps

1. Set organization virtual datacenter allocation model to **Reservation Pool**.
2. Configure the organization administrator role to preserve ExtraConfig parameters.

4.2 Configure vCloud Director to Assign Unique UUIDs

By default vCloud Director configuration, VMs created from a vApp template are assigned the same UUID. For vMRF deployments, it must be ensured that newly created VMs are assigned unique UUIDs.

Prerequisites

- The vCloud Director installation directory is known.

Steps

1. Use the following command to check vCloud Director settings for assigning UUIDs for VMs:

```
<vCloud_Director_Installation_Directory>/bin/cell-management-tool  
manage-config --lookup --name backend.cloneBiosUuidOnVmCopy
```

```
Property "backend.cloneBiosUuidOnVmCopy" has value "1".
```



- 1 in the command printout means that VMs are created with the same UUID. The default value is 1.
- Set the value of the `CloneBiosUuidOnVmCopy` parameter to 0 with the following command to ensure that each VM is created with a unique UUID:

```
<vCloud_Director_Installation_Directory>/bin/cell-management-tool  
manage-config -n backend.cloneBiosUuidOnVmCopy -v 0
```
 - Restart all vCloud Director cells in the server group.

4.3 Download and Extract vMRF Software Delivery Package

Before the deployment, the end user must download and extract the vMRF software delivery package. Both the end user and the cloud administrator must have access to the proper example files in the package.

Steps

- Download the vMRF software delivery package to a computer from which the cloud service clients are reachable.
- Extract the vMRF software delivery package.
- Check that the following files exist after extracting the vMRF software delivery package:

— Deployment files (ovf files)

OVF File Name	Description
<code>vmrf.ovf</code>	OVF file for deployment with automatic IP address allocation from IP pools.
<code>vmrf_man_ip.ovf</code>	OVF file for deployment with manual IP address allocation.

— vMRF image (vmdk file)



5 vMRF Deployment Preparations for the Cloud Administrator

The procedures for vMRF deployment preparation must be performed by the cloud administrator to prepare the cloud environment for running vMRF. The procedures described in this section serve as examples only to demonstrate how to fulfill the vMRF requirements.

5.1 Cloud Hardware and Software Preparation and Configuration

Preparation for vMRF deployment starts by checking that the necessary hardware exists, and making hardware-related configurations in the VIM, in the hypervisor, and in the host Operating System.

5.1.1 Group Compute Nodes for vMRF into DRS Cluster

Perform this procedure only if you want to specify exactly which hosts can run vMRF due to, for example hardware considerations.

Steps

1. Create a Distributed Resource Scheduler (DRS) cluster for the hosts selected to run vMRF. For the details, see [Using DRS Clusters to Manage Resources](#) in the VMware documentation.
2. Use VM-Host affinity rules to define a relationship between vMRF VMs and the group of hosts selected to run vMRF. For the details, see [VM-Host Affinity Rules](#) in the VMware documentation.

5.2 Create Network Topology

The vMRF VNF instance connects to networks. The networks in [Table 1](#) must be created already before the VNF instance can be deployed, since the OVF template uses them as input parameters.

Table 1 vMRF Networks

Network Type
VNF-internal ⁽¹⁾
O&M
H.248 signaling towards MTAS



Network Type
User plane towards media networks

(1) Each VNF instance version requires a dedicated VNF-internal network.

Steps

1. Using the network plan, create the required networks listed in [Table 1](#) depending on vNIC configuration, the following connectivity, if they do not exist.

Depending on the IP allocation alternative, the following additional configuration options apply:

- `vmrf.ovf`: Organization Networks must be configured with IP pool, subnet, subnet mask length, and gateway, and associated to the networks in [Table 1](#).

Note: Fake IP subnets for signaling and media can be used for trusted and untrusted networks. vMRF does not use these IP addresses, but the used IP addresses are configured in the vMRF MOM.

- `vmrf_man_ip.ovf`: Organization Networks must be configured with subnet, subnet mask length, and gateway, and associated to the networks in [Table 1](#).
2. Inform the personnel who are doing the deployment.



6 vMRF Deployment for the End User

After the deployment preparations are completed by the cloud administrator, the end user can start vMRF deployment.

6.1 Initial VNF Configuration Data for Deployment

Initial configuration data means the data needed for vMRF to start processing traffic. This procedure describes how to prepare initial configuration data if you are importing it during deployment. For other options, see [Initial Configuration Guide](#).

While providing cloud configuration data during deployment is mandatory, importing initial configuration is optional. It can be provided during deployment in Base64 encoding. The input can be generated with the following command:

```
base64 -w 0 mrsvconfig.tar.gz
```

If the size of the initial configuration data file `mrsvconfig.tar.gz` exceeds the 23 kB limit, it can be imported in an `iso` file. For more information, see [Create VNF Configuration ISO File](#) on page 10.

6.1.1 Create VNF Configuration ISO File

Initial configuration data must be imported to the VNF in an `iso` if it exceeds the limit imposed by vSphere and cannot be provided during deployment with cloud configuration data.

Steps

1. Rename the `exported_config.tar.gz` file to `mrsvconfig.tar.gz`.
2. Create an `iso` file which includes the `mrsvconfig.tar.gz` file. For example, on a Linux computer, use the following command:

```
genisoimage -l -iso-level 4 -o mrs-init.iso mrsvconfig.tar.gz
```



7 Configure Shared Storage

Note: This procedure is optional. Perform the steps only if shared storage is used.

The external shared storage allows for the storage of the following files from each VM on a remote server:

- Log files from the `/var/log` directory, including journal log files
- Crash dump files
- Configuration backup files created by the automatic backup and restore function

vMRF connects to the server with SSHFS, mounts a specified directory path, and creates a subfolder for the cluster, and subfolders for the files for each VM in the cluster. This ensures that logs and other shared files of different VMs and VNF instances do not get mixed up.

For authentication, an SSH key pair has to be created. This key pair can be cluster-specific, or common for all clusters.

The following parameters must be prepared and included in the `example_environment.yaml` file parameter list or the vApp property list during deployment:

- Storage server username
- Storage server address (IP address and port number)
- Storage server path
- Storage server SSH private key and fingerprint

Prerequisites

- The private and public SSH keys are generated.

Steps

1. Generate the private SSH key parameter value by replacing the end of line characters with the string `"\n"` and including the key data string between single quotation mark (') characters using the following command:

```
echo "'$(awk 'BEGIN{ORS="\n"} {print $0}' .ssh/id_rsa)'"
```



Note: The resulting string must be added in OpenStack-based deployments as the `shared_storage_ssh_private_key` value in the `example_environment.yaml` file or, in VMware-based deployments as vApp property.

2. Generate the SSH fingerprint parameter value from the public SSH key by replacing the end of line characters with the string `"\n"` using the following command:

```
echo "'$(ssh-keyscan -p <server_port> <server_host> | awk  
'BEGIN{ORS="\n"} {print $0}')"'
```

Note: The resulting string must be added in OpenStack-based deployments as the `shared_storage_server_fingerprint_key` value in the `example_environment.yaml` file or, in VMware-based deployments as vApp property.

3. On the shared storage server, append the public key to the authorized keys file using the following command:

```
cat .ssh/<public_key_file_name> >> .ssh/<authorized_keys_file_name>
```

— RELATED INFORMATION —

[9. Deploy vApp from the Catalog on page 14](#)



8 Onboard to the Catalog

This procedure describes the steps needed for onboarding the OVF file to the Catalog.

Steps

1. On the **vApp Templates** tab, click the **Upload...** icon, and select the OVF file to upload the package to the Catalog.
2. If necessary, upload the initial configuration iso file to the catalog, under **Media & Other**.

Note: The image has to contain the configuration with the name `mrsvconfig.tar.gz`.



9 Deploy vApp from the Catalog

This procedure describes how to deploy the vApp from the Catalog.

1. Select the vApp template from **All Templates**, and click **Next**.
2. Name the vApp, define its location, select a proper vCD, and click **Next**.
3. Name the VM within the vApp, and click **Next**.
4. Map network resources to their proper destinations. Set the IP allocation method according to the selected OVF package, and click **Next**.

OVF File Name	Description
vmrf.ovf	OVF file for deployment with automatic IP address allocation from IP pools.
vmrf_man_ip.ovf	OVF file for deployment with manual IP address allocation.

5. Give values to parameters and click **Next**.

Note: Mandatory parameters are marked by a red frame.

The following restrictions apply when defining IP address and password parameter values:

- IP addresses must be in the O&M network IP range.
- Passwords must be generated with the following command:
mkpasswd -m sha-512 <Password>

For platform automatic IP address allocation, provide the following vApp properties:

vApp Property Type	vApp Property
Emergency user credentials	Username
	Password hash. The default password must be changed by supplying a new password hash in this parameter. ⁽¹⁾
	Ssh public key
O&M IP address of the VNF	Movable IP address , an IPv4 address in dot-decimal notation
NTP server IP address	NTP IPs , that is, the list of NTP server addresses separated by a space character
Shared storage configuration ⁽²⁾	Shared storage server username
	Shared storage server path ⁽³⁾
	Shared storage server IP
	Shared storage server port



vApp Property Type	vApp Property
	Shared storage server fingerprint ⁽⁴⁾
	Shared storage ssh private key ⁽⁵⁾
Announcement storage configuration ⁽²⁾	Announcement storage server username
	Announcement storage server path ⁽⁶⁾
	Announcement storage server IP
	Announcement storage server port
	Announcement shared storage server fingerprint ⁽⁴⁾
	Announcement storage ssh private key ⁽⁵⁾
PM data monitoring configuration ⁽²⁾	Pm data monitoring hosts IP address
	Pm data monitoring hosts port
Initial configuration ⁽²⁾	Initial configuration

(1) Password change for the emergency user is supported only during deployment.

(2) Optional property type.

(3) The external shared storage POSIX path pointing to a directory on the shared storage server.

Note: vMRF creates VM specific folders under this path. In order to separate VMs of different vApps, include vApp identifier in this path.

(4) The input can be generated by replacing the end of line characters with the string "\n" and including the key data string between single quotation mark (') characters with the following command: `echo "$(ssh-keyscan -p <server_port> <server_ip_address> | awk 'BEGIN{ORS="\n"} {print $0}')"`.

(5) The input can be generated by replacing the end of line characters with the string "\n" and including the key data string between single quotation mark (') characters using the following command: `echo "$(awk 'BEGIN{ORS="\n"} {print $0}' .ssh/id_rsa)"`.

(6) The announcement shared storage POSIX path pointing to a directory on the shared storage server.

6. Customize VM resources according to your needs and in line with minimum requirements described in vMRF Infrastructure Requirements, and click **Next**.
7. Review the vApp configurations.
8. If the optional initial configuration is properly set up, or it is not provided during deployment, select the **Power on vApp after this wizard is finished** option. Otherwise continue with the next step.
9. Click **Finish**.
10. For the optional steps regarding providing initial configuration, continue with the following procedure, otherwise move on to [Power on vMRF vApp](#) on page 18.

If the **Power on vApp after this wizard is finished** was not selected in [Step 8](#), and initial configuration is to be provided during deployment.



11. Continue with [Provide Initial Configuration from VNF Configuration ISO File](#) on page 17.



10 Provide Initial Configuration from VNF Configuration ISO File

This procedure describes how to provide initial configuration from an iso file created in [Create VNF Configuration ISO File](#) on page 10.

Note: This procedure is optional. Only do the following steps if configuration is needed from an iso file.

Steps

1. Right-click on the VM.
2. Select **Insert CD/DVD** from Catalog.
3. Select the proper iso file.
4. Click **Insert**.



11 Power on vMRF vApp

This procedure describes how to power on the vMRF vApp that is now in the inventory. After powering on the vMRF vApp, if initial configuration was provided during deployment time, vMRF starts running traffic.

Note: If no configuration was provided before power on, the VM starts processing traffic only after configuration is completed either manually, or by importing it.

Steps

1. Navigate to the VM in the vApp.
2. Right-click the VM, and select **Power On**.

If a delay is set in the startup settings of the VM, the VM is powered on only after the set time expires.

3. Check if the VM works correctly as described in [Check vMRF Status](#) on page 20.



12 Scaling Out to Full VNF Size

These procedures describe how to scale out to the full size of the VNF.

12.1 Scale-out Using the Catalog

This procedure describes how to scale out using the Catalog.

Steps

1. Click on the **vApp** name on the **My Cloud > vApps** tab.
2. Select **Add VM** to add VMs.
 - a. Select the VM in the Catalog.
 - b. Click **Add** to add new VMs.
 - c. Click **Next**.
3. Map networks to vNICs based on the settings of the existing VM, select IP address allocation method, and click **Next**.

Note: If manual IP address allocation is used, provide IP addresses for the networks, and make sure that these IP addresses are not colliding with IP addresses configured in VM **Guest Properties** after scale-out but before power-on.

IPv4, IPv6, or both addresses can be assigned to the user plane network towards media networks.

IPv4 addresses must be assigned to the other networks.

The networks must be mapped to vNICs according to the following table:

vNIC	Network Type
vNIC0	VNF-internal
vNIC1	O&M
vNIC2	Signaling
vNIC3	User plane towards media networks

4. Power on the VM.



13 Check vMRF Status

This procedure describes how to verify the vMRF deployment. The status check involves running a vMRF command.

Steps

1. Open an SSH connection to the O&M IP address of the vMRF VNF instance using the following command:

```
ssh -A <user_ID>@<O&M_IP_address>
```

2. Run the following command to check the status of VMs in the cluster:

```
verify_vmrf_cluster_status.py
```

- If the VMs are operating correctly, the OK status is displayed in the command printout. You can exit this procedure.
- If there are VMs with faulty components, a list of faulty VMs and detailed component information of the cluster is displayed. Continue with the next step.

3. Check all components with erroneous state. For specific trouble cases and remedies, refer to the [vMRF Troubleshooting Guideline](#).

Note: The `MrfDirector` and `COM` components are in the OK state only for the VM whose `SC role` is `ACTIVE`, that is, the active System Controller (SC) VM. In all other VMs, these components are in the `OK, NOT RUNNING` state, which is normal behavior.