

Oracle® SD-WAN Edge

High Availability Guide



Release 8.2
F25038-02
June 2021

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

Copyright © 2014, 2021, Oracle and/or its affiliates.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

About This Guide

My Oracle Support iv

Revision History

1 High Availability

Configuring High Availability 1-1
Selecting a High Availability Mode 1-4
Summary 1-7

2 High Availability Configuration for Virtual Appliances

Configuring High Availability on KVM 2-1
 Create Linux/Networking Bridge 2-1
Configuring High Availability on VMWare ESXi 2-3

About This Guide

The purpose of this document is to describe how to implement Oracle Talari Appliance High Availability (HA), as well as various deployments and configurations.

Documentation Set

The following table lists related documentation.

| Document Name | Document Description |
|--|---|
| Oracle SD-WAN Edge Release Notes | Contains information about added features, resolved issues, requirements for use, and known issues in the latest Oracle SD-WAN Edge release. |
| Oracle SD-WAN OS Release Notes and Upgrade Guide | Contains information about inserting an OS Partition Image or OS Patch on an appliance in order to migrate to a new OS version or apply fixes to an existing version. |
| Oracle SD-WAN Security Guide | Contains information about security methods within the Oracle SD-WAN solution. |
| Oracle SD-WAN Edge Features Guide | Contains feature descriptions and procedures for all incremental releases of Oracle SD-WAN Edge. This guide is organized by release version. |

My Oracle Support

My Oracle Support (<https://support.oracle.com>) is your initial point of contact for all product support and training needs. A representative at Customer Access Support (CAS) can assist you with My Oracle Support registration.

Call the CAS main number at 1-800-223-1711 (toll-free in the US), or call the Oracle Support hotline for your local country from the list at <http://www.oracle.com/us/support/contact/index.html>. When calling, make the selections in the sequence shown below on the Support telephone menu:

1. Select 2 for New Service Request.
2. Select 3 for Hardware, Networking, and Solaris Operating System Support.
3. Select one of the following options:
 - For technical issues such as creating a new Service Request (SR), select 1.
 - For non-technical issues such as registration or assistance with My Oracle Support, select 2.

You are connected to a live agent who can assist you with My Oracle Support registration and opening a support ticket.

My Oracle Support is available 24 hours a day, 7 days a week, 365 days a year.

Emergency Response

In the event of a critical service situation, emergency response is offered by the Customer Access Support (CAS) main number at 1-800-223-1711 (toll-free in the US), or call the Oracle Support hotline for your local country from the list at <http://www.oracle.com/us/support/contact/index.html>. The emergency response provides immediate coverage, automatic escalation, and other features to ensure that the critical situation is resolved as rapidly as possible.

A critical situation is defined as a problem with the installed equipment that severely affects service, traffic, or maintenance capabilities, and requires immediate corrective action. Critical situations affect service and/or system operation resulting in one or several of these situations:

- A total system failure that results in loss of all transaction processing capability
- Significant reduction in system capacity or traffic handling capability
- Loss of the system's ability to perform automatic system reconfiguration
- Inability to restart a processor or the system
- Corruption of system databases that requires service affecting corrective actions
- Loss of access for maintenance or recovery operations
- Loss of the system ability to provide any required critical or major trouble notification

Any other problem severely affecting service, capacity/traffic, billing, and maintenance capabilities may be defined as critical by prior discussion and agreement with Oracle.

Locate Product Documentation on the Oracle Help Center Site

Oracle Communications customer documentation is available on the web at the Oracle Help Center (OHC) site, <http://docs.oracle.com>. You do not have to register to access these documents. Viewing these files requires Adobe Acrobat Reader, which can be downloaded at <http://www.adobe.com>.

1. Access the Oracle Help Center site at <http://docs.oracle.com>.
2. Click **Industries**.
3. Click the **Oracle Communications** link.
Under the **SD-WAN** header, select a product.
4. Select the Release Number.
A list of the entire documentation set for the selected product and release appears.
5. To download a file to your location, right-click the **PDF** link, select **Save target as** (or similar command based on your browser), and save to a local folder.

Revision History

This section provides a revision history for this document.

| Date | Description |
|---------------|---|
| February 2020 | <ul style="list-style-type: none">Initial release of this publication, including 8.2M1 feature "High Availability Configuration for Virtual Appliances" |
| June 2021 | <ul style="list-style-type: none">Adds guidelines for configuring MAC addresses in the Configuring HA Properties section in "Configuring High Availability" |

1

High Availability

Appliances can be deployed in High Availability (HA) configuration as a pair of appliances in Active/Standby roles. There are three modes of HA deployment:

- Parallel Inline HA
- Serial Inline HA
- One Arm HA

These HA deployment modes are similar to Virtual Router Redundancy Protocol (VRRP) but use a proprietary protocol called Redundant APN Control Protocol (RACP). Both Client Nodes (Clients) and Network Control Nodes (NCNs) within a Oracle Adaptive Private Network (APN) can be deployed in an HA configuration, if the selected Appliance model supports HA. The T510 and E50 do not support HA; all other appliance models do.

Note:

The NCN is the central Appliance that acts as the master controller of the APN, as well as the central point of administration for the Clients. The NCNs primary purpose is to establish and utilize Conduits with one or more Clients across the network for enterprise Site-to-Site communications.

In HA configuration, one Appliance at the Site is designated the Active appliance and is continuously monitored by the Standby appliance. Configuration is mirrored across both appliances. If the Standby appliance loses connectivity with the Active one for a defined period of time, the Standby appliance assumes the identity of the Active appliance and takes over the traffic load. Depending on the deployment mode this fast failover has minimal impact on the application traffic flowing through the Site. We will discuss the impact in more detail later in this document.

Note: For NCNs, we also support what is called Geographically-Diverse NCN redundancy. In this mode, one of the Clients is also designated as a secondary NCN. It will continuously monitor the health of the Primary NCN and if a catastrophic event occurs, it will assume the role of the NCN. The T510, T730, and E50 appliance models cannot act as NCNs

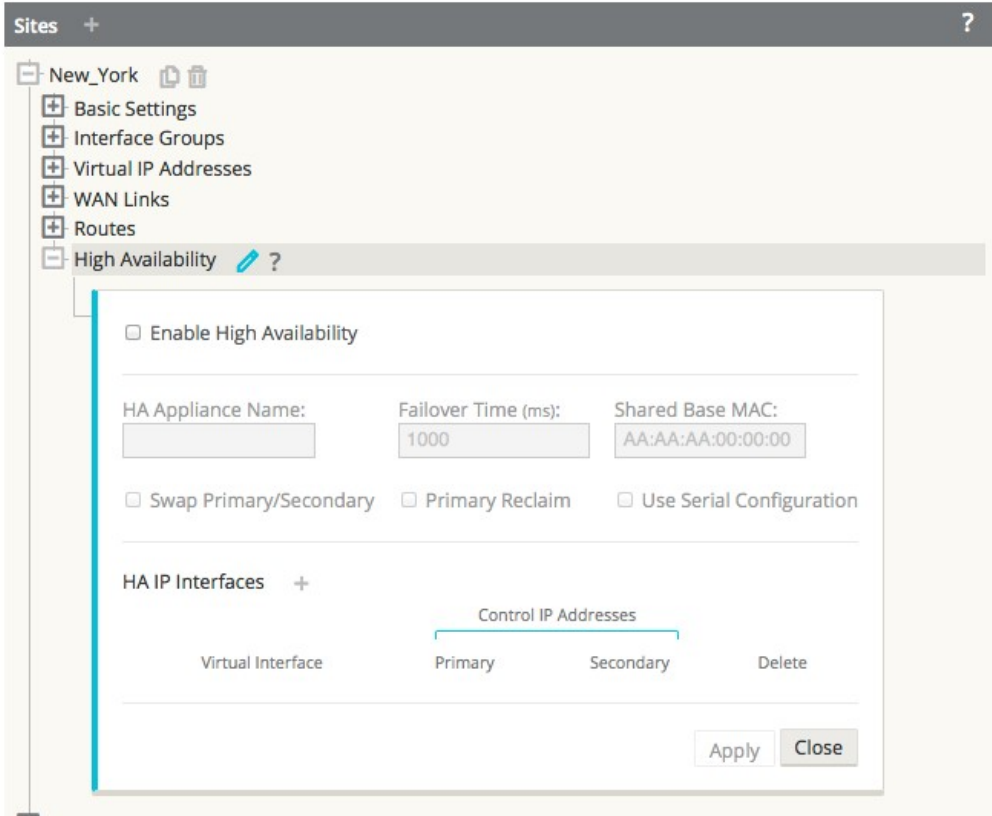
There are various technical considerations in each deployment scenario. These will be explored in the sections below.

Configuring High Availability

See the following sections for configuring HA.

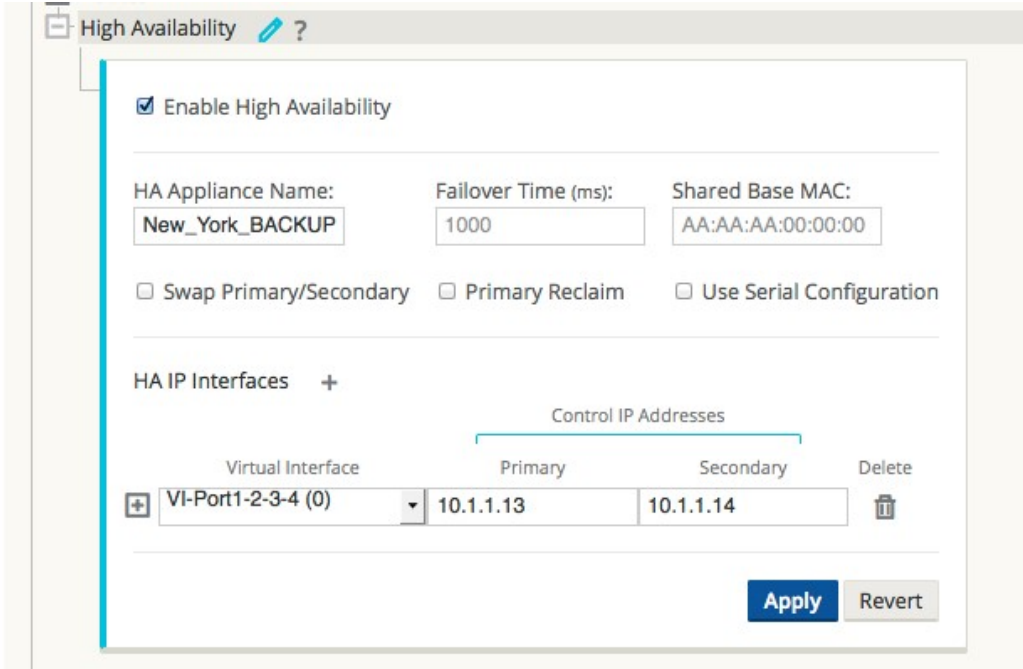
Configuring a Site for HA

HA is configured through the Oracle Configuration Editor tool. When a Site is added, an HA appliance can be configured for the Site:



Configuring HA Properties

Once a Site has been configured with an HA appliance, the HA appliance and interface groups can be configured:



 **Note:**

For sites configured with HA, all interfaces communicating with devices other than its HA peer will use a virtual MAC address. You can configure the virtual MAC address by providing a base MAC for the site in the fields shown in the image above. A virtual MAC will then be created for each interface at the site by adding the interface number to the least significant byte of the base MAC address.

In a virtual environment, if multiple sites are configured with HA, you need to configure different base MAC. However, you must ensure the difference is not in the least significant byte. For example, for site 2, use aa:aa:aa:00:02:00. For site 3, use aa:aa:aa:00:03:00.

Primary Reclaim

In the event that the Active appliance fails and then comes back, it can be configured to reclaim the Active status once it has rebooted. This feature is disabled by default. To enable it, select the check box for “Primary Reclaim” in the High Availability section of the site configuration.





The Active/Standby states of an HA pair can be manually switched from the web console of either appliance during run-time operation.


Serial Inline HA


When Serial Inline HA mode is desired, select the check box for “Use Serial Configuration.”

Interface Groups

At least one HA interface group must be configured. This is the interface that the HA RACP protocol will be established across in order to monitor the Active appliance for reachability. For One Arm HA mode, only one interface group is required. For Inline HA mode, additional interface groups may be configured in order to use External Tracking to monitor reachability of the upstream or downstream network infrastructure (e.g. switch port failure) to detect if an HA state change is needed.

| | Virtual Interface | Primary | Secondary | Delete |
|---|--------------------|-----------|-----------|---|
|  | | | |  |
|  | VI-Port1-2-3-4 (0) | 10.1.1.13 | 10.1.1.14 |  |

External Tracking 

| External Tracker IP Address | Interface | Delete |
|-----------------------------|-----------|---|
| 10.1.1.1 | 2 |  |

Selecting a High Availability Mode


Select a mode for high availability.

One Arm HA

In One Arm mode, the HA appliance pair is outside of the data path. Application traffic of interest is redirected to the appliance pair, typically using Policy Based Routing (PBR). One Arm is used when a single insertion point in the network is not feasible or to avoid the challenges of fail-to-wire.

In this case, adding HA is straight forward. The Standby appliance can simply be added to the same VLAN or subnet as the Active appliance and the router, as we show in the diagram below:

VLAN or subnet as the Active appliance and the router, as we show in the diagram below:



In One Arm mode it is recommended that the Talari Appliances do not reside in the data network subnets. This means the Talari Conduit traffic doesn't have to traverse the PBR and avoids route loops,

In One Arm mode it is recommended that the Appliances do not reside in the data network subnets. This means the Oracle Conduit traffic doesn't have to traverse the PBR and avoids route loops, etc. The Oracle Appliances and router do have to be directly connected, either via an Ethernet port or by residing in the same VLAN.

Using IP SLA Monitoring for Fall Back


As long as one of the Appliances is active, traffic will still flow even if the Conduit is down. In this case, the Appliance will redirect the traffic back to the router as Intranet traffic. However, if both Appliances become disabled, the router will still try to redirect traffic to the appliances. IP SLA monitoring can be configured at the router to disable the PBR if the next device is not reachable. This allows the router to fall back to doing a route lookup in the normal way and forwarding packets appropriately.

 **Note:**

Not all routers and firewalls support PBR or IP SLA.

Parallel Inline HA

In Parallel Inline HA mode, the Appliances are deployed alongside each other, in line with the data path. The diagram below shows a common deployment with multiple switches and a single router.




In the above diagram, only one path through the Active appliance is used. It is important to note the bypass interface groups are configured to be fail-to-block and not fail-to-wire so that we don't get spanning tree loops during a failover.

The HA state can be monitored through the inline interface groups or through a direct connection between the appliances. External Tracking can be used to monitor the reachability of the upstream or downstream network infrastructure (e.g. switch port failure) to detect if an HA state change is needed. If both Appliances are disabled or fail, a tertiary path can be used directly between the switch and router. This path must have a higher spanning tree cost than the Appliance paths so that it is not used under normal conditions.

Failover in Parallel Inline HA mode is very quick and nearly hitless, as no physical state change occurs. Fallback to the tertiary path is not typically hitless and can cause traffic to be blocked for 5-30 seconds depending on the spanning tree configuration.

If there are out of path connections to other WAN Links, both appliances must be connected to them. In more complex scenarios, where multiple routers might be using VRRP, non-

routable VLANs are recommended to ensure the LAN side switches and WAN side routers are reachable at Layer 2.



Serial Inline HA

In Serial Inline HA mode, the Appliances are inline on the same path. In this case the bypass interface groups should be in the fail-to-wire mode, with the Standby appliance in a Passthrough or bypass state. A direct connection between the two appliances on a separate port must be configured and used for the HA interface group. Serial Inline has the advantage of being very simple to deploy but has some drawbacks:

- Due to a physical state change when the Appliance switches over from Active to Standby, failover can cause some loss of connectivity depending on how long the auto-negotiation takes on the Ethernet ports. It is likely to be several seconds and can be service impacting.
- It is not recommended that Serial Inline be used on ports that are auto-negotiated, as this will increase failover time.
- If the HA connection between the appliances fails in some way, both appliances will go active and cause a service interruption. This can be mitigated by assigning multiple HA connections so there is no single point of failure.
- We recommend testing fully when inline with other devices, using the following scenarios to verify bypass (fail-to-wire) operation.
 - Appliance In-Line: **Powered OFF**
 - Appliance In-Line: **Powered ON with Talari Service DISABLED**
 - Appliance In-Line: **Powered ON with Talari Service ENABLED**

An example of Serial Inline HA deployment is shown below:

HA-INTERFACE CONNECTION

Summary

The three modes of HA deployment and their advantages and disadvantages are summarized in the table below:

| Deployment Mode | Configuration Complexity | Physical Complexity | Failover Time | Fallback |
|-----------------|--------------------------|---------------------|---------------|-------------------------|
| One Arm | High (PBR) | Low | Fast <1s | Yes (Intranet & IP SLA) |
| Parallel Inline | Medium | Medium (VLANS) | Fast <1s | Yes (High Cost Path) |
| Serial Inline | Medium | Low | Slow 5-15s | Yes (Passthrough) |

As a rule of thumb, either One Arm HA configuration or Parallel Inline HA configuration is recommended for Datacenters or Sites that forward a high volume of traffic to minimize disruption during failover. If a small loss of service is acceptable during a failover, then Serial Inline is a reasonable solution.

Serial Inline HA protects against appliance failure and Parallel Inline HA protects against all failures. In all cases, HA is valuable to preserve the continuity of the APN during a system failure.

2

High Availability Configuration for Virtual Appliances

Linux KVM and ESXi can be deployed using HA configuration. This allows one Active appliance to be monitored by a Standby appliance. In case of failover, the Standby appliance mirrors the configuration of the Active appliance and overtakes the traffic load.



Note:

Serial HA is not supported for Virtual Appliances.

Configuring High Availability on KVM

To support HA, an instance should be created using Linux bridge on the KVM server.

Create Linux/Networking Bridge

Follow these instructions to create a networking bridge.

1. Log in to the KVM server.
2. Create a file called `ifcfg-lanbrN` and replace N with the interface number under `/etc/sysconfig/network-scripts/`.
3. Open the file in an editor and enter the following

```
[localadmin@localhost network-scripts]$ cat ifcfg-lanbr201
DEVICE=lanbr201
TYPE=Bridge
BOOTPROTO=none
ONBOOT=yes
DELAY=0
[localadmin@localhost network-scripts]$
```

4. To add the virtual interface to the LAN bridge, ensure `ONBOOT=yes` and `BRIDGE=` the name of the LAN bridge in the `ifcfg-ens2f0` file, where `ifcfg-ens2f0` is the virtual interface.

```
[localadmin@localhost network-scripts]$ cat ifcfg-ens2f0
TYPE=Ethernet
PROXY_METHOD=none
BROWSER_ONLY=no
BOOTPROTO=dhcp
DEFROUTE=yes
IPV4_FAILURE_FATAL=no
IPV6INIT=yes
```

```
IPV6_AUTOCONF=yes
IPV6_DEFROUTE=yes
IPV6_FAILURE_FATAL=no
IPV6_ADDR_GEN_MODE=stable-privacy
NAME=ens2f0
UUID=bf4196e3-b003-41ff-8b02-29ed79ea3552
DEVICE=ens2f0
ONBOOT=yes
BRIDGE=lanbr201
[localadmin@localhost network-scripts]$
```

5. Create a WAN bridge by logging into the KVM server.
6. Create a file called `ifcfg-wanbrN` and replace N with the interface number under `/etc/sysconfig/network-scripts`.
7. Open the file in an editor and enter the following.

```
[localadmin@localhost network-scripts]$ cat ifcfg-wanbr201
DEVICE=wanbr201
TYPE=Bridge
BOOTPROTO=none
ONBOOT=yes
DELAY=0
[localadmin@localhost network-scripts]$
```

8. To add the virtual interface to the WAN bridge, ensure `ONBOOT=yes` and `BRIDGE=`the name of the WAN bridge in the `ifcfg-ens2f1` file, where `ifcfg-ens2f1` is the virtual interface.

```
[localadmin@localhost network-scripts]$ cat ifcfg-ens2f1
TYPE=Ethernet
PROXY_METHOD=none
BROWSER_ONLY=no
BOOTPROTO=dhcp
DEFROUTE=yes
IPV4_FAILURE_FATAL=no
IPV6INIT=yes
IPV6_AUTOCONF=yes
IPV6_DEFROUTE=yes
IPV6_FAILURE_FATAL=no
IPV6_ADDR_GEN_MODE=stable-privacy
NAME=ens2f1
UUID=f45577ab-f733-4c53-a791-fe44662cc5b4
DEVICE=ens2f1
ONBOOT=yes
BRIDGE=wanbr201
[localadmin@localhost network-scripts]$
```

9. Restart the network by entering the following.

```
$sudo systemctl restart network
```

10. Verify the interfaces are connected to the bridges by entering the following.

```
$sudo brctl show
```

The interfaces should look like the following

```
[localadmin@localhost network-scripts]$ brctl show
bridge name      bridge id        STP enabled     interfaces
lanbr201         8000.3cfdfe6272a8  no              ens2f0
                                                         vnet0
lanbr202         8000.3cfdfe6272aa  no              ens2f2
                                                         vnet1
lanbr203         8000.3cfdfe6272b8  no              ens3f0
                                                         vnet2

wanbr201         8000.3cfdfe6272a9  no              ens2f1
                                                         vnet3
wanbr202         8000.3cfdfe6272ab  no              ens2f3
                                                         vnet4
wanbr203         8000.3cfdfe6272b9  no              ens3f1
                                                         vnet5
                                                         vnet6

[localadmin@localhost network-scripts]$
```

Configuring High Availability on VMWare ESXi

For instructions on how to set up HA sites and configure HA properties, see "Configuring High Availability."

1. After deploying a VM, create a standard vSwitch.
2. Remove the Uplink1 by clicking on the X button.
3. Create a new port group and add the vSwitch you created in Step 1.
4. Attach Network Adapter 5 of the Active appliance to the vSwitch.
5. Attach Network Adapter 5 of the Standby appliance to the vSwitch.

Note:

Make sure the Network Adapter attached to the Standby appliance matches the one attached to the Active appliance.

6. From the editor tool, enable the Oracle SD-WAN Edge service.
7. Create an HA interface group with the following configuration
 - Ethernet Interfaces: 4
 - Bypass mode: Fail-to-Block
8. Assign the virtual IP to the HA interface.