High Availability Overview

The goal of any high availability solution is to ensure that all network services are resilient to failure. Such a solution aims to provide continuous access to network resources by addressing the potential causes of downtime through functionality, design, and best practices. The core of the Cisco Catalyst SD-WAN high availability solution is achieved through a combination of three factors:

- Functional hardware device redundancy. The basic strategy consists of installing and provisioning redundant hardware devices and redundant components on the hardware. These devices are connected by a secure control plane mesh of Datagram Transport Layer Security (DTLS) connections among themselves, which allows for rapid failover should a device fail or otherwise become unavailable. A key feature of the Cisco Catalyst SD-WAN control plane is that it is established and maintained automatically, by the Cisco vEdge devices and software themselves.
- Robust network design.
- Software mechanisms ensure rapid recovery from a failure. To provide a resilient control plane, the Cisco Catalyst SD-WAN Overlay Management Protocol (OMP) regularly monitors the status of all Cisco vEdge devices in the network and automatically adjusts to changes in the topology as devices join and leave the network. For data plane resiliency, the Cisco Catalyst SD-WAN software implements standard protocol mechanisms, specifically Bidirectional Forwarding Detection (BFD), which runs on the secure IPsec tunnels between routers.

Recovery from a failure is a function of the time it takes to detect the failure and then repair or recover from it. TheCisco Catalyst SD-WAN solution provides the ability to control the amount of time to detect a failure in the network. In most cases, repair of the failure is fairly instantaneous.

Hardware Support of High Availability

A standard best practice in any network setup is to install redundant hardware at all levels, including duplicate parallel routers and other systems, redundant fans, power supplies and other hardware components within these devices, and backup network connections. Providing high availability in the Cisco Catalyst SD-WAN solution is no different. A network design that is resilient in the face of hardware failure should include redundant Cisco Catalyst SD-WAN Validators, Cisco Catalyst SD-WAN Controllers, and routers and any available redundant hardware components.

Recovery from the total failure of a hardware component in the Cisco Catalyst SD-WAN overlay network happens in basically the same way as in any other network. A backup component has been preconfigured, and it is able to perform all necessary functions by itself.

Robust Network Design

In addition to simple duplication of hardware components, the high availability of a Cisco Catalyst SD-WAN network can be enhanced by following best practices to design a network that is robust in the face of failure. In one such network design, redundant components are spread around the network as much as possible. Design practices include situating redundant Cisco Catalyst SD-WAN Validators and Cisco Catalyst SD-WAN Controllers at dispersed geographical locations and connecting them to different transport networks. Similarly, the routers at a local site can connect to different transport networks and can reach these networks through different NATs and DMZs.

Software Support of High Availability

The Cisco Catalyst SD-WAN software support for high availability and resiliency in the face of failure is provided both in the control plane, using the standard DTLS protocol and the proprietary Cisco Catalyst SD-WAN Overlay Management Protocol (OMP), and in the data plane, using the industry-standard protocols BFD, BGP, OSPF, and VRRP.

Control Plane Software Support of High Availability

The Cisco Catalyst SD-WAN control plane operates in conjunction with redundant components to ensure that the overlay network remains resilient if one of the components fails. The control plane is built on top of DTLS connections between the Cisco devices, and it is monitored by the Cisco Catalyst SD-WAN OMP protocol, which establishes peering sessions (similar to BGP peering sessions) between pairs of vSmart controllers and routers, and between pairs of vSmart controllers. These peering sessions allow OMP to monitor the status of the Cisco devices and to share the information among them so that each device in the network has a consistent view of the overlay network. The exchange of control plane information over OMP peering sessions is a key piece in the Cisco Catalyst SD-WAN high availability solution:

- vSmart controllers quickly and automatically learn when a vBond orchestrator or a router joins or leaves the network. They can then rapidly make the necessary modifications in the route information that they send to the routers.
- vBond orchestrators quickly and automatically learn when a device joins the network and when a vSmart controller leaves the network. They can then rapidly make the necessary changes to the list of vSmart controller IP addresses that they send to routers joining the network.
- vBond orchestrators learn when a domain has multiple vSmart controllers and can then provide multiple vSmart controller addresses to routers joining the network.
- vSmart controllers learn about the presence of other vSmart controllers, and they all automatically synchronize their route tables. If one vSmart controller fails, the remaining systems take over management of the control plane, simply and automatically, and all routers in the network continue to receive current, consistent routing and TLOC updates from the remaining vSmart controllers.

Let's look at the redundancy provided by each of the Cisco Catalyst SD-WAN hardware devices in support of network high availability.

Recovering from a Failure in the Control Plane

The combination of hardware component redundancy with the architecture of the Cisco Catalyst SD-WAN control plane results in a highly available network, one that continues to operate normally and without interruption when a failure occurs in one of the redundant control plane components. Recovery from the total failure of a vSmart controller, vBond orchestrator, or router in the Cisco Catalyst SD-WAN overlay network

happens in basically the same way as the recovery from the failure of a regular router or server on the network: A preconfigured backup component is able to perform all necessary functions by itself.

In the Cisco Catalyst SD-WAN solution, when a network device fails and a redundant device is present, network operation continues without interruption. This is true for all Cisco devices—vBond orchestrators, vSmart controllers, and routers. No user configuration is required to implement this behavior; it happens automatically. The OMP peering sessions running between Cisco devices ensure that all the devices have a current and accurate view of the network topology.

Let's examine failure recovery device by device.

Data Plane Software Support for High Availability

For data plane resiliency, the Cisco Catalyst SD-WAN software implements the standard BFD protocol, which runs automatically on the secure IPsec connections between routers. These IPsec connections are used for the data plane, and for data traffic, and are independent of the DTLS tunnels used by the control plane. BFD is used to detect connection failures between the routers. It measures data loss and latency on the data tunnel to determine the status of the devices at either end of the connection.

BFD is enabled, by default, on connections between Cisco IOS XE Catalyst SD-WAN devices and Cisco vEdge devices. BFD sends Hello packets periodically (by default, every 1 second) to determine whether the session is still operational. If a certain number of the Hello packets are not received, BFD considers that the link has failed and brings the BFD session down (the default dead time is 3 seconds). When BFD sessions goes down, any route that points to a next hop over that IPsec tunnel is removed from the forwarding table (FIB), but it is still present in the route table (RIB).

In the Cisco Catalyst SD-WAN software, you can adjust the Hello packet and dead time intervals. If the timers on the two ends of a BFD link are different, BFD negotiates to use the lower value.

Using Affinity To Manage Network Scaling

In the Cisco Catalyst SD-WAN overlay network, all Cisco IOS XE Catalyst SD-WAN devices and Cisco vEdge devices establish control connections to all vSmart controllers, to ensure that the routers are always able to properly route data traffic across the network. As networks increase in size, with routers at thousands of sites and with vSmart controllers in multiple data centers managing the flow of control and data traffic among routers, network operation can be improved by limiting the number of vSmart controllers that a router can connect to. When data centers are distributed across a broad geography, network operation can also be better managed by having routers establish control connections only with the vSmart controllers collocated in the same geographic region.

Establishing affinity between vSmart controllers and Cisco vEdge devices allow you to control the scaling of the overlay network, by limiting the number of vSmart controllers that a Cisco vEdge device can establish control connections (and form TLOCs) with. When you have redundant routers in a single data center, affinity allows you to distribute the vEdge control connections across the vSmart controllers. Similarly, when you have multiple data centers in the overlay network, affinity allows you to distribute the vEdge control connections across the data centers. With affinity, you can also define primary and backup control connections, to maintain overlay network operation in case the connection to a single vSmart controller or to a single data center fails.

A simple case for using affinity is when redundant vSmart controllers are collocated in a single data center. Together, these vSmart controllers service all the Cisco vEdge devices in the overlay network. The figure below illustrates this situation, showing a scenario with three vSmart controllers in the data center and, for simplicity, showing just three of the many Cisco vEdge devices in the network.

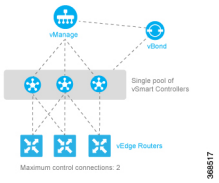

If you do not enable affinity, each Cisco vEdge device establishes a control connection—that is, a TLOC—to each of the three vSmart controllers in the data center. Thus, a total of nine TLOCs are established, and each router exchanges OMP updates with each controller. Having this many TLOCs can strain the resources of both the vSmart controllers and the Cisco vEdge devices , and the strain increases in networks with larger numbers of Cisco vEdge devices .

Enabling affinity allows each Cisco vEdge device to establish TLOC connections with only a subset of the vSmart controllers. The figure above shows each router connecting to just two of the three vSmart controllers, thus reducing the total number of TLOCs from nine to six. Both TLOC connections can be active, for a total of six control connections. It is also possible for one of the TLOC connections be the primary, or preferred, and the other to be a backup, to be used as an alternate only when the primary is unavailable, thus reducing the number of active TLOCs to three.

Affinity also enables redundancy among data centers, for a scenario in which multiple vSmart controllers are collocated in two or more data centers. Then, if the link between a Cisco vEdge device and one of the data centers goes down, the vSmart controllers in the second data center are available to continue servicing the overlay network. The figure below illustrates this scenario, showing three vSmart controllers in each of two data centers. Each of the three Cisco vEdge devices establishes a TLOC connection to one controller in the West data center and one in the East data center.

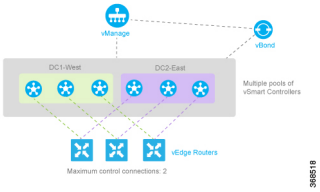

You might think of the scenario in the figure above as one where there are redundant data centers in the same region of the world, such as in the same city, province, or country. For an overlay network that spans a larger geography, such as across a continent or across multiple continents, you can use affinity to limit the network scale either by restricting Cisco vEdge devices so that they connect only to local vSmart controllers or by having Cisco vEdge devices preferentially establish control connections with data centers that are in their geographic region. With geographic affinity, Cisco vEdge devices establish their only or their primary TLOC connection or connections with vSmart controllers in more local data centers, but they have a backup available to a more distant region to provide redundancy in case the closer data centers become unavailable. The figure below illustrates this scenario, Here, the Cisco vEdge devices in Europe have their primary TLOC connections to the two European data centers and alternate connections to the data centers in North America. Similarly, for the Cisco vEdge devices in North America, the primary connections are to the two North American data centers, and the backup connections are to the two European data centers.

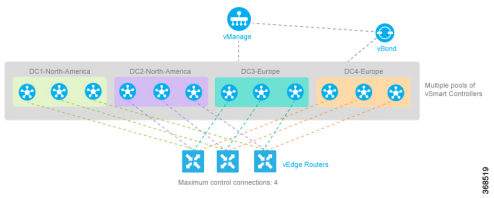

As is the case with any overlay network that has multiple vSmart controllers, all policy configurations on all the vSmart controllers must be the same.

Before you cofigure High availability, to start with the configuration transaction, you can use the following command such as,

```
ntp server 198.51.241.229 source GigabitEthernet1 version 4
```


Cisco Catalyst SD-WAN Validator Redundancy

The Cisco SD-WAN Validator performs two key functions in the Cisco Catalyst SD-WAN overlay network:

- Authenticates and validates all Cisco SD-WAN Controllers and routers that attempt to join the Cisco Catalyst SD-WAN network.
- Orchestrates the control plane connections between the Cisco SD-WAN Controllers and routers, thus enabling Cisco SD-WAN Controller and routers to connect to each other in the Cisco Catalyst SD-WAN network.

The Cisco SD-WAN Validator runs as a VM on a network server. The Cisco SD-WAN Validator can also run on a router that is configured to be a Cisco SD-WAN Validator, however this is not recommended, and it limits the number of router control connections to 50. If using running the Cisco SD-WAN Validator daemon on a router, note that only one Cisco SD-WAN Validator daemon can run at a time on a router, so to provide redundancy and high availability, the network must have two or more routers that function as Cisco SD-WAN Validator orchestrators. (Note also that it is not recommended to use a router acting as a vBond orchestrator as a regular router.)

Having multiple Cisco SD-WAN Validators ensures that one of them is always available whenever a Cisco device such as a router or a Cisco SD-WAN Controller is attempting to join the network.

Configuration of Redundant Cisco Catalyst SD-WAN Validators

A Cisco SD-WAN Controller learns that it is acting as a Cisco SD-WAN Validator from its configuration. In the **system vbond** configuration command, which defines the IP address (or addresses) of the Cisco SD-WAN Validator (or validators) in the Cisco Catalyst SD-WAN overlay network, you include the **local** option. In this command, you also include the local public IP address of the Cisco SD-WAN Validator, (Even though on Cisco vEdge device and Cisco SD-WAN Controllers you can specify an IP address of Cisco SD-WAN Validator as a DNS name, on the Cisco SD-WAN Validator itself, you must specify it as an IP address.)

On Cisco SD-WAN Controllers, and Cisco vEdge devices, when the network has only a single Cisco SD-WAN Validator, you can configure the location of the Cisco SD-WAN Validator system either as an IP address or as the name of a DNS server (such as vbond.cisco.com). (Again, you configure this in the **system vbond** command.) When the network has two or more Cisco SD-WAN Validators and they must all be reachable, you should use the name of a DNS server. The DNS server then resolves the name to a single IP address that the Cisco SD-WAN Validator returns to the Cisco vEdge device. If the DNS name resolves to multiple IP addresses, the Cisco SD-WAN Validator returns them all to the Cisco vEdge device, and the router tries each address sequentially until it forms a successful connection.

Note that even if your Cisco Catalyst SD-WAN network has only a single Cisco SD-WAN Validator, it is recommended as a best practice that you specify a DNS name rather than an IP address in the **system vbond** configuration command, because this results in a scalable configuration. Then, if you add additional Cisco SD-WAN Validators to your network, you do not need to change the configurations on any of the routers or Cisco SD-WAN Controllers in your network.

Recovering from a Cisco Catalyst SD-WAN Validator Failure

In a network with multiple Cisco SD-WAN Validators, if one of them fails, the other Cisco SD-WAN Validators simply continue operating and are able to handle all requests by Cisco devices to join the network. From a control plane point of view, each Cisco SD-WAN Validator maintains a permanent DTLS connections to each of the Cisco SD-WAN Controllers in the network. (Note however, that there are no connections between the Cisco SD-WAN Validators themselves.) As long as one Cisco SD-WAN Validator is present in the domain, the Cisco Catalyst SD-WAN network is able to continue operating without interruption, because Cisco SD-WAN Controllers and routers are still able to locate each other and join the network.

Because Cisco SD-WAN Validators never participate in the data plane of the overlay network, the failure of any Cisco SD-WAN Validator has no impact on data traffic. Cisco SD-WAN Validators communicate with routers only when the routers are first joining the network. The joining router establishes a transient DTLS connection with a Cisco SD-WAN Validator to learn the IP address of a Cisco SD-WAN Controller. When the Cisco vEdge device configuration lists the Cisco SD-WAN Validator address as a DNS name, the router tries each of the Cisco SD-WAN Validators in the list, one by one, until it is able to establish a DTLS connection. This mechanism allows a router to always be able to join the network, even after one of a group of Cisco SD-WAN Validators has failed.

Cisco Catalyst SD-WAN Manager Server Redundancy

The Cisco SD-WAN Manager servers comprise a centralized network management system that enables configuration and management of the Cisco devices in the overlay network. It also provides a real-time dashboard into the status of the network and network devices. The Cisco SD-WAN Manager servers maintain permanent communication channels with all Cisco vEdge devices in the network. Over these channels, the Cisco SD-WAN Manager servers push out files that list the serial numbers of all valid devices, they push out the configuration of each device, and they push out new software images as part of a software upgrade process. From each network device, the Cisco SD-WAN Manager servers receive various status information that is displayed on the Cisco SD-WAN Manager **Monitor** > **Overview** page.

Note In Cisco vManage Release 20.6.1 and earlier releases, the status information is available on the **Dashboard** > **Main Dashboard** page.

A highly available Cisco Catalyst SD-WAN network contains three or more Cisco SD-WAN Manager servers in each domain. This scenario is referred to as a cluster of Cisco SD-WAN Manager servers, and each Cisco SD-WAN Manager server in a cluster is referred to as a Cisco SD-WAN Manager instance. Each Cisco SD-WAN Manager instance in a cluster can manage approximately 2000 devices, so a cluster of three Cisco SD-WAN Manager instances can manage up to 6000 devices. The Cisco SD-WAN Manager instances automatically load balances the devices that they manage. With three instances, the Cisco SD-WAN Manager cluster remains operational if one of the devices in that cluster fail.

For related information, see Troubleshooting TechNotes.

A Cisco SD-WAN Manager cluster consists of the following architectural components:

- Application server—This provides a web server for user sessions. Through these sessions, a logged-in user can view a high-level dashboard summary of networks events and status, and can drill down to view details of these events. A user can also manage network serial number files, certificates, software upgrades, device reboots, and configuration of the Cisco SD-WAN Manager cluster itself from the Cisco SD-WAN Manager application server.
- Configuration database—Stores the inventory and state and the configurations for all Cisco vEdge devices.
- Network configuration system—Stores all configuration information, policies, templates, certificates, and more.
- Statistics database—Stores the statistics information collected from all Cisco devices in the overlay network.
- Message bus—Communication bus among the different Cisco SD-WAN Manager instances. This bus is used to share data and coordinate operations among the Cisco SD-WAN Manager instances in the cluster.

The Statistics database and Configuration database services must run on an odd number of Cisco SD-WAN Manager instances, with a minimum of three. For these databases to be writeable, there must be a quorum of Cisco SD-WAN Manager instances running and they should be in sync. A quorum is a simple majority. For example, if you have a cluster of three Cisco SD-WAN Manager instances running these databases, then two must be running and in sync. Initially, all Cisco SD-WAN Manager instances run the same services. However, you can choose not to run some services on some instances. From the **Cluster Management** window, you can select the services that can run on each Cisco SD-WAN Manager instance. You can add a fourth Cisco SD-WAN Manager instance to load balance more Cisco vEdge device s. In such a case, disable the statistics database and configuration database on one of the instances Cisco SD-WAN Manager because those services need to run on an odd number of instances. Optionally, you can run the configuration database on a single instance to reduce the amount of information shared between the devices and reduce load.

The following figure shows the interaction between Cisco SD-WAN Manager instances in a cluster, although a minimum of three devices are required. The figure illustrates the Cisco SD-WAN Manager services that synchronize between the Cisco SD-WAN Manager instances. Also in this figure, you see that each Cisco SD-WAN Manager instance resides on a virtual machine (VM). The VM can have from one to eight cores, with a Cisco Catalyst SD-WAN software process (vdaemon) running on each core. In addition, the VM stores the actual configuration for the Cisco SD-WAN Manager server itself.

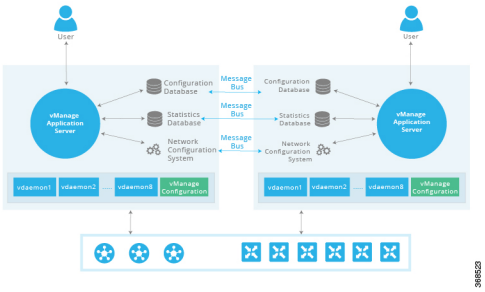

The Cisco SD-WAN Manager cluster implements an active-active architecture in the following way:

- Each of the Cisco SD-WAN Manager instance in the cluster is an independent processing node.
- All Cisco SD-WAN Manager instances are active simultaneously.
- All user sessions to the application server are load balanced by using an external load balancer.

- All control sessions between the Cisco SD-WAN Manager application servers and the routers are load balanced. A single Cisco SD-WAN Manager instance can manage a maximum of about 2000 Cisco vEdge devices . However, all the controller sessions—the sessions between the Cisco SD-WAN Manager instances and the Cisco Catalyst SD-WAN Controllers, the sessions between the Cisco SD-WAN Manager instances and the Cisco Catalyst SD-WAN Validators are arranged in a full-mesh topology.
- The configuration and statistics databases can be replicated across Cisco SD-WAN Manager instances, and these databases can be accessed and used by all the Cisco SD-WAN Manager instances.
- If one of the Cisco SD-WAN Manager instances in the cluster fails or otherwise becomes unavailable, the network management services that are provided by the Cisco SD-WAN Manager server are still fully available across the network.

The message bus among the Cisco SD-WAN Manager instances in the cluster allows all the instances to communicate using an out-of-band network. This design, which leverages a third vNIC on the Cisco SD-WAN Manager VM, avoids using WAN bandwidth for management traffic.

You configure the Cisco SD-WAN Manager cluster from the Cisco SD-WAN Manager web application server. During the configuration process, you can configure each Cisco SD-WAN Manager instance that can run the following services:

- Application server—Each Cisco SD-WAN Manager server runs an application server instance.
- Configuration database—Within the Cisco SD-WAN Manager cluster, no more than three iterations of the configuration database can run.
- Load balancer—The Cisco SD-WAN Manager cluster requires a load balancer to distribute user login sessions among the Cisco SD-WAN Manager instances in the cluster. As mentioned, a single Cisco SD-WAN Manager instance can manage a maximum of about 2000 WAN edge devices.
- Messaging server—We recommend that you configure each Cisco SD-WAN Manager instance to run the message bus so that all the instances in the cluster can communicate with each other.
- Statistics database—Within a Cisco SD-WAN Manager cluster, no more than three iterations of the statistics database can run.
- Coordination server: It’s used internally by the Messaging server.

The following are the design considerations for a Cisco SD-WAN Manager cluster:

- A Cisco SD-WAN Manager cluster should consist of a minimum of three Cisco SD-WAN Manager instances.
- The application server and message bus should run on all Cisco SD-WAN Manager instances.
- Within a cluster, a maximum of three instances of the configuration database and three instances of the statistics database can run. However, any individual Cisco SD-WAN Manager instance can run both, one, or none of these two databases.
- To provide the greatest availability, we recommend that you run the configuration and statistics databases on three Cisco SD-WAN Manager instances.

Deploy Cisco Catalyst SD-WAN Manager Cluster

Note **Prerequisites for Cisco SD-WAN Releases 20.3.1, 20.3.2, 20.3.2.1, and 20.4.1 Only**

Starting with Cisco SD-WAN Release 20.3.1, you must run the messaging server on all the active instances of the Cisco SD-WAN Manager cluster, and form a full-mesh messaging cluster.

When upgrading to Cisco SD-WAN Releases 20.3.1 and later from earlier releases, ensure that you run the messaging server on each of the Cisco SD-WAN Manager instance, prior to the upgrade.

Ensure that you have a minimum of three Cisco SD-WAN Manager instances in a Cisco SD-WAN Manager cluster.

The deployment process requires multiple reboots and should be performed during a scheduled maintenance window.

1. Back up the Cisco SD-WAN Manager database. Use the following command to back up.

   request nms configuration-db backup path /home/admin/<db _ backup _ filename>

2. If the current device has only two network interface cards (NICs), add a third NIC, don't use Dynamic Host Configuration Protocol (DHCP) for addressing. This third NIC is used for cluster messaging between the devices ofCisco SD-WAN Manager instances, within vpn 0. For a device to detect the new interface, the device must be rebooted. Configure the interface and verify that it has connectivity.

3. From the Cisco SD-WAN Manager menu, choose **Administration** > **Cluster Management**, edit the IP address to change localhost to the IP address of the third NIC, which is used for cluster messaging.

4. Restart the Cisco SD-WAN Manager services. This may take some time. You can view the */var/log/nms/vmanage-server.log* for the log output to stabilize, and then use the **request nms all status** command to determine for how long the processes have been running. When it comes up, verify that Cisco SD-WAN Manager is operational and stable.

5. Provision two more Cisco SD-WAN Manager VMs in your virtual environment with the appropriate disk size and third NIC.

6. Configure two more Cisco SD-WAN Manager VMs with minimal system configuration and addressing for the NICs. Configure the admin user password to match that is configured on the original device. If you're using enterprise certificates, ensure that you install the root certificate chain on the new device as you did with the first device. Also, ensure that the clocks of the new devices are in sync with the original device.

   The following is a sample of a minimal configuration:

```
system
            host-name             vManage3
            system-ip             10.0.1.102
            site-id               1
            organization-name     cisco-sdwan1
            vbond vbond!
            vpn 0
                host vbond ip 198.51.100.103 192.51.100.104
                interface eth0
                ip address 198.51.100.102/24
                tunnel-interface
                no shutdown
                !
                interface eth1
```

```
ip address 198.51.100.102/24
no shutdown !
ip route 0.0.0.0/0 0.0.0.0
!
vpn 512
interface eth2
    ip address 198.56.101.102/24
    no shutdown
!
```

Note While a default gateway is given for the out of band Cisco SD-WAN Manager cluster interfaces, it's not required if the Cisco SD-WAN Manager cluster nodes are using addresses in the same subnet and are adjacent.

7. From the Cisco SD-WAN Manager menu, choose **Administration** > **Cluster Management**, to add one of the new Cisco SD-WAN Manager VMs to the cluster by adding the IP address of the third NIC for database replication, click **Add vManage**. Choose all services.

   The processes for the second Cisco SD-WAN Manager instance restarts. This process might take some time. View the */var/log/nms/vmanage-server.log* and then use the **request nms all status** command to determine process completion time.

8. View the new device on the second Cisco SD-WAN Manager instances, by clicking **Configuration** > **Certificates** > **Controllers**.

9. Generate a certificate signing request (CSR), get the device signed, and install the signed device certificate for this new device. For more information, see Cluster Management.

   The cluster shows that the new device on the second instance of Cisco SD-WAN Manager is rebalancing and that the Cisco SD-WAN Manager services are restarting on the device on previous Cisco SD-WAN Manager instance. A new task appears in the task bar for the **Cluster Sync**. Although the task appears as Complete, view the */var/log/nms/vmanage-server.log* to resolve errors, if any.

   The first instance of Cisco SD-WAN Manager also restarts services, which eventually reset the GUI session. It might take several minutes for the GUI to become available after the services restarts. View the */var/log/nms/vmanage-server.log*, and then use the **request nms all status** command.

10. Wait for **Cluster Sync** to complete. View the *vmanage-server.log* to resolve errors, if any.

11. After the Cluster Sync is completed, add the second of the new Cisco SD-WAN Manager VMs to the cluster by adding the IP address of the third NIC for database replication. Select all services.

    The third instance of Cisco SD-WAN Manager restarts services. This might take some time. A new task appears in the task bar for the **Cluster Sync**. The first and second instances of Cisco SD-WAN Manager also restart services, which eventually reset the GUI session. It may take several minutes for the GUI to become available after the Cisco SD-WAN Manager services restart. The device on the new Cisco SD-WAN Manager instance appears on the **Configuration** > **Certificates** > **Controllers** page. Perform steps 9 and 10 for this device.

Upgrade Cisco Catalyst SD-WAN Manager Cluster

To upgrade a cluster, ensure that the services start in an orderly manner. After the upgrade steps, use the steps in the Restarting the NMS Processes section to start all the services in an orderly manner.

To get the software partitions prepared to be activated, upgrade the devices (without activating).

1. Take backup of the Cisco SD-WAN Manager server. This can take a while. If Cisco hosts the controllers and there's a recent snapshot, this step can be skipped.

   ```
   request nms configuration-db backup path/home/admin/<db _ backup _ filename>
   ```

2. Stop Cisco SD-WAN Manager services on all devices in the cluster by using the following command on each device.

   ```
   request nms all stop
   ```

3. Activate the new version on each device. This activation causes each device to reload.

   ```
   request software activate <version>
   ```

4. If you do the upgrade from CLI, ensure that you manually confirm the upgrade from the CLI after the reload and before it reverts to the previous version.

   ```
   request software upgrade-confirm
   ```

5. After the devices reboot, stop Cisco SD-WAN Manager services on all devices in the cluster.

   ```
   request nms all stop
   ```

Next, ensure that you perform the steps of restarting the Cisco SD-WAN Manager server manually.

Restarting the Cisco Catalyst SD-WAN Manager Server Manually

When the cluster is in a bad state as part of the upgrade, you should manually restart the Cisco SD-WAN Manager server. To start the Cisco SD-WAN Manager server, restart the processes one at a time in an orderly manner instead of using **request nms all restart** or a similar command. The following manual restart order might vary for your cluster, depending on what services you're running on the devices in the cluster. The following order is based on a basic cluster with three Cisco SD-WAN Manager instances.

Note Consider bringing up the services manually as mentioned in the following method whenever you have to reboot a Cisco SD-WAN Manager server or after an upgrade.

1. On each device, stop all Cisco SD-WAN Manager services.

   ```
   request nms all stop
   ```

2. Verify that all services have stopped. It's normal for the preceding command to give some message about failing to stop a service if it takes too long, so use the following command to verify that everything is stopped before proceeding.

   ```
   request nms all status
   ```

3. Start the Statistics database on each device that is configured to run it. Wait for the service to start each time before proceeding to the next device.

   ```
   request nms statistics-db start
   ```

4. Verify that the service is started before proceeding to start it on the next device. After service starts, perform step 3 to start the Statistics database on the next device. Once all the devices, ensure that the Statistics database is running, and then proceed to the next step.

   ```
   request nms statistics-db status
   ```

5. Start the Configuration database on each device that is configured to run it. Wait for the service to start each time before proceeding to the next device.

```
request nms configuration-db start
```

6. Verify that the service has started before proceeding to start it on the next device. Go to vshell and tail a log file to look for a database is online message. When confirmed, go to step 5 to start the Configuration database on the next device. After all devices have the Configuration database running, proceed to the next step.

   ```
   tail -f -n 100 /var/log/nms/debug.log
   ```

Note For versions prior to 18.2, use this log file

```
tail -f -n 100 /var/log/nms/vmanage-orientdb-database.log
```

7. Start the Coordination server on each device. Wait for the service to start each time before proceeding to the next device.

   ```
   request nms coordination-server start
   ```

8. Verify that the service is started before proceeding to start it on the next device. After verifying, go to step 7 to start the Coordination server on the next device. After the Coordination server runs on all the devices, proceed to the next step.

   ```
   request nms coordination-server status
   ```

9. Start the Messaging server on each device. Wait for the service to start each time before proceeding to the next device.

   ```
   request nms messaging-server start
   ```

10. Verify that the service has started before proceeding to start it on the next device. After verifying, go to step 9 to start the Messaging server on the next device. After the Messaging server runs on all devices, proceed to the next step.

    ```
    request nms messaging-server status
    ```

11. Start the Application server on each device. Wait for the service to start each time before proceeding to the next device.

    ```
    request nms application-server start
    ```

12. Verify that the service has started before proceeding to start it on the next device. To verify if the service is fully started, open the GUI of that device. After the GUI is fully loaded and you're able to log in, go to step 11 to start the Application server on the next device.

13. Restart the Cisco SD-WAN Manager server cloud services on each device. Wait for the service to start each time before proceeding to the next device.

    ```
     request nms cloud-agent start
    ```

14. Verify that the service has started before proceeding to start it on the next device. After verifying, go to step 12 to start the cloud services on the next device. After the cloud services run on all devices, proceed to the next step.

    ```
    request nms cloud-agent status
    ```

15. To verify that there are no errors and everything has loaded cleanly, tail the log files.

Check Cisco SD-WAN Manager to verify that all devices appear as online and reachable, and that the statistics exist.

Cisco Catalyst SD-WAN Manager Backups

Cisco manages Cisco SD-WAN Manager by taking regular snapshots of the devices for recovery due to a catastrophic failure or corruption. The frequency and retention of these snapshots are set for each overlay. Generally, the snapshots are taken daily and retained for up to 10 days. For certain scheduled maintenance activities, such as the upgrade of the devices, another snapshot can be taken before the scheduled activity. In all other cases, it's your responsibility to take regular backups of the Cisco SD-WAN Manager configuration database and snapshots of the Cisco SD-WAN Manager virtual machine, and follow the example of frequency and retention that is followed by Cisco.

Cisco Catalyst SD-WAN Manager Database Backup

Although the Cisco SD-WAN Manager cluster provides high availability and a level of fault tolerance, regular backup of the configuration database should be taken and stored securely off-site. Cisco SD-WAN Manager doesn't have a mechanism of automating the collection of a configuration database backup on a schedule and copying it to another server. The greater the time between the backup and when it's needed for a recovery, the greater the risk that data might be lost. Perform configuration database backups often. Use the following command to create a configuration database backup file.

```
request nms configuration-db backup path <path>
```

Cisco Catalyst SD-WAN Controller Redundancy

Cisco Catalyst SD-WAN Controller Redundancy

The Cisco SD-WAN Controllers are the central orchestrators of the control plane. They have permanent communication channels with all the Cisco devices in the network. Over the DTLS connections between the Cisco SD-WAN Controllers and Cisco SD-WAN Validators and between pairs of Cisco SD-WAN Controllers, the devices regularly exchange their views of the network, to ensure that their route tables remain synchronized. The Cisco SD-WAN Controllers pass accurate and timely route information over DTLS connections to Cisco vEdge device .

A highly available Cisco Catalyst SD-WAN network contains two or more Cisco SD-WAN Controllers in each domain. A Cisco Catalyst SD-WAN domain can have up to 20 Cisco SD-WAN Controllers, and each router, by default, connects to two of them. When the number of Cisco SD-WAN Controllers in a domain is greater than the maximum number of controllers that a domain's routers are allowed to connect to, the Cisco Catalyst SD-WAN software load-balances the connections among the available Cisco SD-WAN Controllers.

While the configurations on all the Cisco SD-WAN Controllers must be functionally similar, the control policies must be identical. This is required to ensure that, at any time, all Cisco vEdge devices receive consistent views of the network. If the control policies are not absolutely identical, different Cisco SD-WAN Controllers might give different information to a Cisco vEdge device , and the likely result will be network connectivity issues.

Note To reiterate, the Cisco Catalyst SD-WAN overlay network works properly only when the control policies on all Cisco SD-WAN Controllers are identical. Even the slightest difference in the policies will result in issues with the functioning of the network.

To remain synchronized with each other, the Cisco SD-WAN Controllers establish a full mesh of DTLS control connections, as well as a full mesh of OMP sessions, between themselves. Over the OMP sessions,

the Cisco SD-WAN Controllers advertise routes, TLOCs, services, policies, and encryption keys. It is this exchange of information that allows the Cisco SD-WAN Controllers to remain synchronized.

You can place Cisco SD-WAN Controllers anywhere in the network. For availability, it is highly recommended that the Cisco SD-WAN Controllers be geographically dispersed.

Each Cisco SD-WAN Controller establishes a permanent DTLS connection to each of the Cisco SD-WAN Validators. These connections allow the Cisco SD-WAN Validators to track which Cisco SD-WAN Controllers are present and operational. So, if one of the Cisco SD-WAN Controller fails, the Cisco SD-WAN Validator does not provide the address of the unavailable Cisco SD-WAN Controller to a router that is just joining the network.

To reiterate, the Cisco Catalyst SD-WAN overlay network works properly only when the control policies on all Cisco SD-WAN Controllers are identical. Even the slightest difference in the policies result in issues with the functioning of the network.

Recovering from a Cisco Catalyst SD-WAN Controller Failure

The Cisco SD-WAN Controllers are the primary controllers of the network. To maintain this control, they maintain permanent DTLS connections to all the Cisco SD-WAN Validators and Cisco IOS XE Catalyst SD-WAN devices and Cisco vEdge devices. These connections allow the Cisco SD-WAN Controllers to be constantly aware of any changes in the network topology. When a network has multiple Cisco SD-WAN Controllers:

- There is a full mesh of OMP sessions among the Cisco SD-WAN Controllers.
- Each Cisco SD-WAN Controller has a permanent DTLS connection to each Cisco SD-WAN Validator.
- The Cisco SD-WAN Controllers have permanent TLS or DTLS connections to the Cisco IOS XE Catalyst SD-WAN devices and Cisco vEdge devices.

If one of the Cisco SD-WAN Controllers fails, the other Cisco SD-WAN Controllers seamlessly take over handling control of the network. The remaining Cisco SD-WAN Controllers are able to work with Cisco IOS XE Catalyst SD-WAN devices and Cisco vEdge devices joining the network and are able to continue sending route updates to the routers. As long as one Cisco SD-WAN Controller is present and operating in the domain, the Cisco Catalyst SD-WAN network may continue operating without interruption.

For more information about configuring graceful restart for OMP, see graceful restart.

Cisco vEdge Device Redundancy

The Cisco vEdge device is commonly used in two ways in the Cisco Catalyst SD-WAN network: to be the Cisco Catalyst SD-WAN routers at a branch site, and to create a hub site that branch routers connect to.

A branch site can have two or more Cisco vEdge device s for redundancy. Each of the router maintains the following connections:

- A secure control plane connection, via a TLS or DTLS connection, with one or more Cisco SD-WAN Controllers in its domain.
- A secure data plane connection with the other routers.

Because both the routers receive the same routing information from the Cisco SD-WAN Controllers, each one is able to continue to route traffic if one fails, even if they are connected to different transport providers.

When using Cisco IOS XE Catalyst SD-WAN devices and Cisco vEdge devices in a hub site, you can provide redundancy by installing two Cisco vEdge device s. The branch routers need to connect to each of the hub routers by using separate DTLS connections.

You can also have Cisco vEdge device provide redundancy by configuring up to tunnel interfaces on a single router. Each tunnel interface can go through the same or different firewalls, service providers, and network clouds, and each maintains a secure control plane connection, by means of a DTLS tunnel, with the Cisco SD-WAN Controllers in its domain.

Recovering from a Cisco vEdge Device Failure

The route tables on Cisco IOS XE Catalyst SD-WAN devices and Cisco vEdge devices are populated by OMP routes received from the Cisco SD-WAN Controllers. For a site or branch with redundant routers, the route tables on both routers remain synchronized, so if either of the routers fail, the other one continues to be able to route data traffic to its destination.

Configure Affinity Between Cisco Catalyst SD-WAN Controllers and Cisco vEdge Device

One way to manage network scale is to configure affinity between Cisco SD-WAN Controllers and Cisco vEdge device s. To do this, you place each Cisco SD-WAN Controller into a controller group, and then you configure which group or groups a Cisco vEdge device can establish control connections with. The controller groups are what establishes the affinity between Cisco SD-WAN Controllers and Cisco vEdge device .

Configure Controller Group Identifier on Cisco Catalyst SD-WAN Controllers

To participate in affinity, each Cisco SD-WAN Controller must be assigned a controller group identifier:

```
vSmart(config)#system controller-group-id number
```

The identifier number can be from 0 through 100.

For Cisco SD-WAN Controllers in the same data center, they can have the same controller group identifier or different identifiers:

- If the Cisco SD-WAN Controllers have the same controller group identifier, a Cisco vEdge device establishes a control connection to any one of them. If that Cisco SD-WAN Controller becomes unreachable, the router simply establishes a control connection with another one of the controllers in the data center. As an example of how this might work, if one Cisco SD-WAN Controller becomes unavailable during a software upgrade, the Cisco vEdge device immediately establishes a new TLOC with another Cisco SD-WAN Controller, and the router's network operation is not interrupted. This network design provides redundancy among Cisco SD-WAN Controllers in a data center.
- If the Cisco SD-WAN Controllers have different controller group identifiers, a Cisco vEdge device can use one controller as the preferred and the other as backup. As an example of how this might work, if you are upgrading the Cisco SD-WAN Controller software, you can upgrade one controller group at a time. If a problem occurs with the upgrade, a Cisco vEdge device establishes TLOCs with the Cisco SD-WAN Controllers in the second, backup controller group, and the router's network operation is not interrupted. When the Cisco SD-WAN Controller in the first group again becomes available, the Cisco vEdge device switches its TLOCs back to that controller. This network design, while offerring redundancy among the Cisco SD-WAN Controllers in a data center, also provides additional fault isolation.

Configure Affinity on Cisco vEdge Device

For a Cisco vEdge device to participate in affinity, you configure the Cisco SD-WAN Controllers that the router is allowed to establish control connections with, and you configure the maximum number of control connections (or TLOCs) that the Cisco vEdge device itself, and that an individual tunnel on the router, is allowed to establish.

Configure Cisco Catalyst SD-WAN Controller Groups

Configuring the Cisco SD-WAN Controllers that the router is allowed to establish control connections is a two-part process:

- At the system level, configure a single list of the controller group identifiers that are present in the overlay network.
- For each tunnel interface, you can choose to restrict which controller group identifiers the tunnel interface can establish control connections with. To do this, configure an exclusion list.

At a system level, configure the identifiers of the Cisco SD-WAN Controller groups:

```
vEdge(config)#system controller-group-list numbers
```

List the Cisco SD-WAN Controller group identifiers that any of the tunnel interfaces on the Cisco vEdge device might want to establish control connections with. It is recommended that this list contain the identifiers for all the Cisco SD-WAN Controller groups in the overlay network.

If you want a specific tunnel interface to establish control connections to only a subset of all the Cisco SD-WAN Controller groups, configure the group identifiers to exclude:

```
vEdge(config-vpn-0-interface)#tunnel-interface exclude-controller-group-list numbers
```

This command lists the identifiers of the Cisco SD-WAN Controller groups that this particular tunnel interface should not establish control connections with, when a Cisco SD-WAN Controller is available in configured controller groups. Ensure that the controller groups in this list are a subset of the controller groups that are configured with the **system controller-group-list** command.

If a Cisco vEdge device is not able to establish control connection with the configured controller group(s), the Cisco vEdge device will try to connect to the vSmart controllers from the excluded list.

To display the controller groups configured on a Cisco vEdge device, use the **show control connections** command.

Configure Maximum Number of Control Connections

By default, the maximum number of control connections that each tunnel interface can establish is the same as the maximum number of OMP sessions that is configured on the Cisco vEdge device. The default value for MOS is 2.

Configuring the maximum number of control connections for a tunnel interface for a Cisco vEdge device is a two-part process:

- At the system level, configure the MOS that the Cisco vEdge device can establish to Cisco SD-WAN Controllers.

- If a tunnel interface needs to connect to a different number Cisco SD-WAN Controllers than the configured MOS value, configure the maximum number of control connections that the tunnel can establish to Cisco SD-WAN Controllers,

Note If MCC is not configured on a tunnel interface, its default value is the same as the MOS value.

Effectively, the maximum number of control connections a tunnel interface in a Cisco vEdge device can establish is determined by the following formula:

Max Control Connections for tunnel interface in VPN 0 = MIN(MOS, MCC)

To modify the maximum number of OMP sessions, enter the following command:

```
vEdge(config)# system max-omp-sessions number
```

A Cisco vEdge device establishes OMP sessions as follows:

- The device, not the individual tunnel interfaces, establishes OMP sessions with Cisco SD-WAN Controllers.
- When different tunnel interfaces on a router connect to the same Cisco SD-WAN Controller, the device creates a single OMP session with the Cisco SD-WAN Controller and the different tunnel interfaces use this single OMP session.

Note When each tunnel interface connects to the same set of Cisco SD-WAN Controllers, a Cisco vEdge device has the total number of OMP sessions equal to the configured maximum number of OMP sessions. However, if each tunnel interface connects to a different Cisco SD-WAN Controller (because of an excluded controller list), the total number of OMP sessions on the device is higher than the configured maximum number of OMP sessions,

Use the following command to modify the maximum number of control connections for a tunnel interface:

```
vEdge(config-vpn-0-interface)# vpn 0 interface interface-name tunnel-interface
max-control-connections number
```

The number of control connections can be from 0 through 100. The default value is the maximum number of OMP sessions that is configured with the **system max-omp-sessions** command.

To display the actual number of control connections for each tunnel interface, use the **show control affinity config** command.

To display a list of the Cisco SD-WAN Controllers that each tunnel interface has established control connections with, use the **show control affinity status** command.

Configure Affinity for Cisco Catalyst SD-WAN Controllers on Single Data Center

In a Cisco Catalyst SD-WAN overlay network that has multiple Cisco SD-WAN Controllers, each tunnel interface in a Cisco vEdge device establishes control connections to two Cisco SD-WAN Controllers. This default behavior provides controller redundancy, so there is no need to configure affinity.

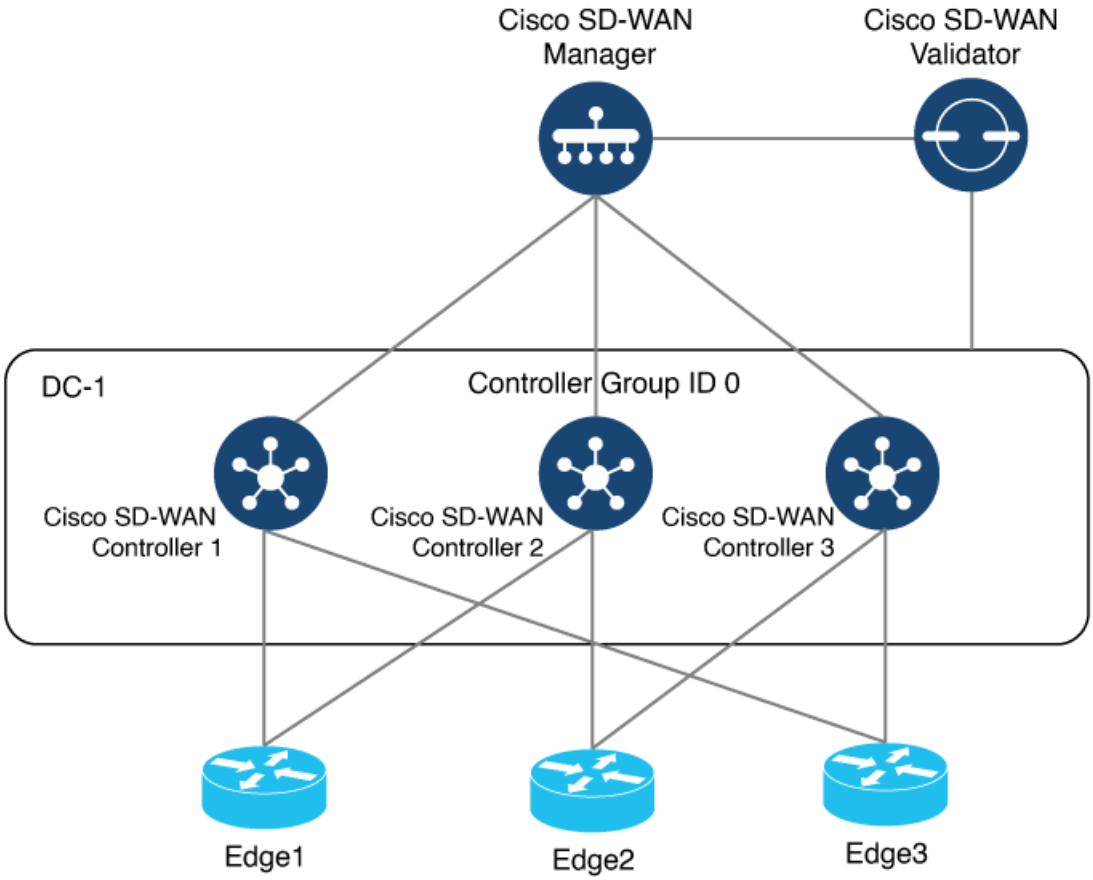

In the topology that is shown in the figure above, there are three Cisco SD-WAN Controllers in Data Center DC-1, all of which belong to default controller group 0. A Cisco vEdge device selects two Cisco SD-WAN Controllers, which are identified as its assigned Cisco SD-WAN Controllers. Each tunnel interface of the Cisco vEdge device connects to these assigned Cisco SD-WAN Controllers. When a tunnel interface connects to its assigned Cisco SD-WAN Controller, that tunnel interface is said to be in *equilibrium*.

Note If exclude-controller-group-list is configured on a tunnel interface, that tunnel interface may have different Cisco SD-WAN Controllers assigned to it.

However, if you want to connect the Cisco vEdge device only to a subset of the Cisco SD-WAN Controllers in a data center, you can use affinity. Place the Cisco SD-WAN Controllers in different controller groups, and then configure the Cisco vEdge device with a list of controller groups that it can connect to. This design provides redundant control connections to the Cisco SD-WAN Controller and provides fault isolation among the Cisco SD-WAN Controller groups in the same data center.

In the topology that is shown in the figure above, assume that you want edge devices to connect to Cisco SD-WAN Controllers as follows:

- Edge1 connects only to Cisco SD-WAN Controller 1 and Cisco SD-WAN Controller 2
- Edge2 connects only to Cisco SD-WAN Controller 2 and Cisco SD-WAN Controller 3

- Edge3 connects only to Cisco SD-WAN Controller 1 and Cisco SD-WAN Controller 3

To achieve these connections, configure Cisco SD-WAN Controller 1 with controller group ID 1:

```
vSmart(config)# system controller-group-id 1
```

To verify the configuration, use the **show running-config** command:

```
vSmart# show running-config system
system
 description         "vSmart in data center 1"
 host-name           vSmart
 gps-location latitude 37.368140
 gps-location longitude -121.913658
 system-ip           172.16.255.19
 site-id             100
 controller-group-id 1
 organization-name   "Cisco"
 clock timezone America/Los_Angeles
```

Use the following commands to configure the other Cisco SD-WAN Controllers:

```
vSmart2(config)# system controller-group-id 2
```

```
vSmart3(config)# system controller-group-id 3
```

In this example, each Cisco SD-WAN Controller in the data center is added to its own controller group. Alternatively, you can add multiple Cisco SD-WAN Controllers to the same controller group.

Note When Cisco SD-WAN Controllers are assigned to controller groups, we recommend that you assign controller groups for all Cisco SD-WAN Controllers in the overlay network.

Next, because you want Edge1 to connect only to Cisco SD-WAN Controller 2 and Cisco SD-WAN Controller 3, configure Edge1 as follows:

```
Edge1(config)# system controller-group-list 2 3
```

Configure Edge2 to connect only to Cisco SD-WAN Controller 1 and Cisco SD-WAN Controller 2:

```
Edge2(config)# system controller-group-list 1 2
```

Configure Edge3 to connect only to Cisco SD-WAN Controller 1 and Cisco SD-WAN Controller 3:

```
Edge3(config)# system controller-group-list 1 3
```

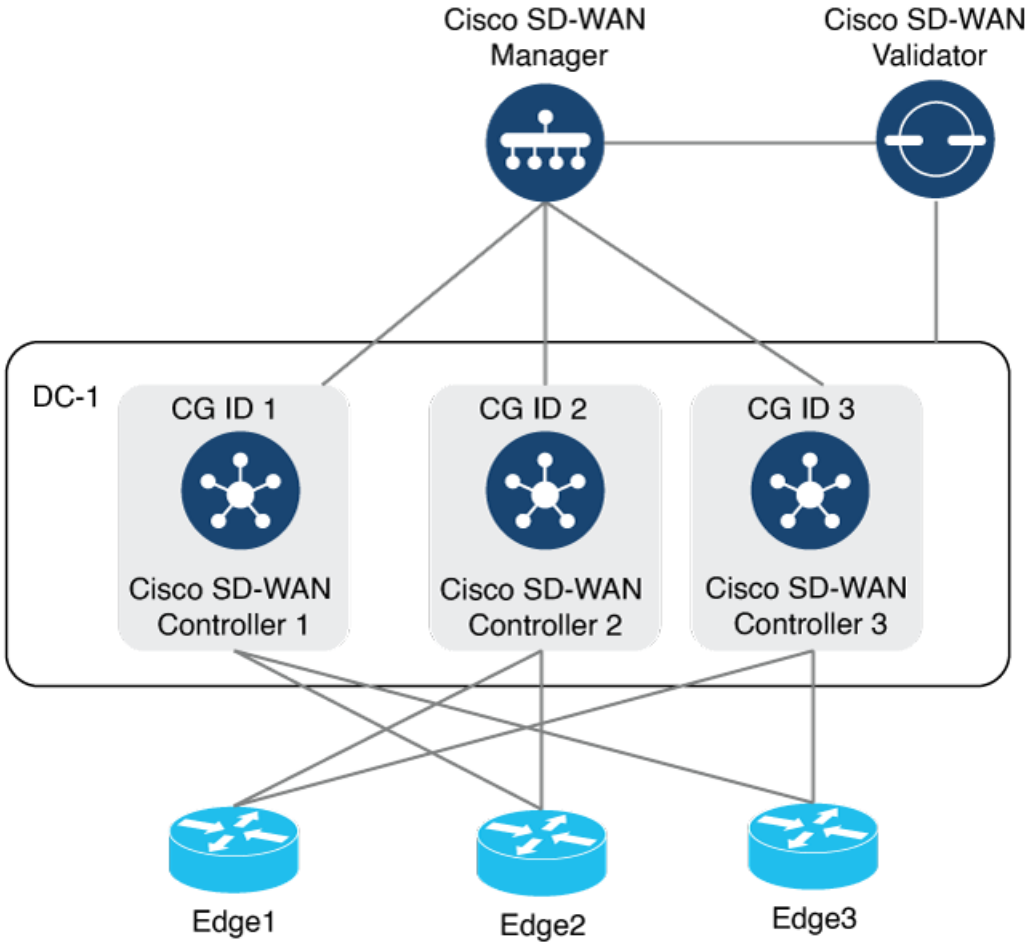

To display the control connections with the Cisco SD-WAN Controllers, use the **show control connections** command. The last column on the output, **Controller Group ID**, lists the Cisco SD-WAN Controller group that a router is in.

```
vEdge-1# show control connections

                                                                    PEER         PEER                                                CONTROLLER
PEER     PEER     PEER             SITE    DOMAIN   PEER          PRIVATE      PEER         PUBLIC                                 GROUP
TYPE     PROTOCOL SYSTEM IP        ID      ID       PRIVATE IP    PORT         PUBLIC IP    PORT    LOCAL COLOR    STATE    UPTIME     ID
-------------------------------------------------------------------------------------------------------------------------------------------
vsmart   dtls     172.16.255.19    100     1        10.0.5.19     12446        10.0.5.19    12446   lte            up       0:00:00:53 1
vsmart   dtls     172.16.255.20    200     1        10.0.12.20    12446        10.0.12.20   12446   lte            up       0:00:00:22 1
```

To display the maximum number of control connections allowed on the router, use the **show control local-properties** command. The last line of the output lists the maximum controllers. The following is the abbreviated output for this command:

```
vEdge-1# show control local-properties

personality                vedge
organization-name          Cisco
certificate-status         Installed
root-ca-chain-status       Installed

certificate-validity       Valid
certificate-not-valid-before Mar 10 19:50:04 2016 GMT
certificate-not-valid-after  Mar 10 19:50:04 2017 GMT
...

                        PUBLIC          PUBLIC PRIVATE        PRIVATE                                PRIVATE                              MAX
```

```
 RESTRICT/           LAST         SPI
TIME    NAT  VM
INTERFACE               IPv4            PORT   IPv4           IPv6                                    PORT    VS/VM COLOR           STATE CNTRL
 CONTROL/     LR/LB  CONNECTION
REMAINING   TYPE CON

 STUN
             PRF
-------------------------------------------------------------------------------------------------------------------------------------------------------------------
------------------
ge0/0                   2.1.1.11        12346  2.1.1.11       ::                                      12346    2/1  default         up     2
   no/yes/no   No/No  2:19:35:21
0:07:58:46  N    5
```

These two commands display information about the control connections established by the affinity configuration. To see, for each interface, which controller groups are configured and which Cisco SD-WAN Controller the interface is connected to, use the **show control affinity config** command:

```
vEdge-1# show control affinity config

EFFECTIVE CONTROLLER LIST FORMAT - G(C),...  - Where G is the Controller Group ID
                                                 C is the Required vSmart Count
CURRENT CONTROLLER LIST FORMAT - G(c)s,...   - Where G is the Controller Group ID
                                                 c is the current vSmart count
                                                 s Status □ when matches, χ when does not match
                 EFFECTIVE
                 REQUIRED
INDEX INTERFACE  VS COUNT   EFFECTIVE CONTROLLER LIST    CURRENT CONTROLLER LIST  EQUILIBRIUM
-----------------------------------------------------------------------------------------------
0     ge0/2      2           1(2)                       1(2)□                    Yes
```

The command output above shows that affinity is configured on interface ge0/2:

Table 1:

Field	Description
Effective Required VS Count	Shows that the interface is configured to create two control connections, and two control connections have been established.
Effective Controller List	Shows that affinity on the interface is configured to use one Cisco SD-WAN Controller from Controller Group 1, shown as 1(1), and one Cisco SD-WAN Controller from Controller Group 2, shown as 2(1). You configure the affinity controller identifiers with the **controller-group-list** command (at the **system** level) and, for the tunnel interface, the **exclude-controller-group-list** command.
Current Controller List	Lists the actual affinity configuration for the interface. The output here shows that the interface has two control connections with Cisco SD-WAN Controller in group 1 and another control connection. The “Y” indicates that the current and effective controller lists match each other.
Equilibrium	Indicates that the current controller lists match what is expected from the affinity configuration for that tunnel interface.

To determine the exact Cisco Catalyst SD-WAN Controllers with which the tunnel interface has established control connections, use the **show control affinity status** command:

```
vEdge-1# show control affinity status

ASSIGNED CONNECTED CONTROLLERS - System IP( G),..       - System IP of the assigned vSmart
                                                             G is the group ID to which the vSmart belongs to

UNASSIGNED CONNECTED CONTROLLERS - System IP( G),..     - System IP of the unassigned vSmart
                                                             G is the group ID to which the vSmart belongs to

INDEX INTERFACE ASSIGNED CONNECTED  CONTROLLERS        UNASSIGNED CONNECTED CONTROLLERS
-------------------------------------------------------------------------------------------------------------
0     ge0/2     172.16.255.19(  1), 172.16.255.20(  1)
```

The command output above shows that interface **ge0/2** has control connections to two Cisco SD-WAN Controllers, 172.16.255.19 and 172.16.255.20, that both controllers are in group 1, and that both controllers are in one of the groups configured in the controller group list. If the interface were connected to a Cisco SD-WAN Controller not in the controller group list, it would be listed in the Unassigned Connected Controllers column.

Configure Affinity for vSmart Controllers on Two Data Centers

You can use affinity to enable redundancy among data centers, for a network design in which multiple Cisco Catalyst SD-WAN Controllers are spread across two or more data centers. Then, if the link between a Cisco vEdge device and one of the data centers goes down, the Cisco Catalyst SD-WAN Controllers in the second data center are available to continue servicing the overlay network. The figure below illustrates this scenario, showing three Cisco Catalyst SD-WAN Controllers in each of two data centers. Each of the three Cisco vEdge device s establishes a TLOC connection to one controller in the West data center and one in the East data center.

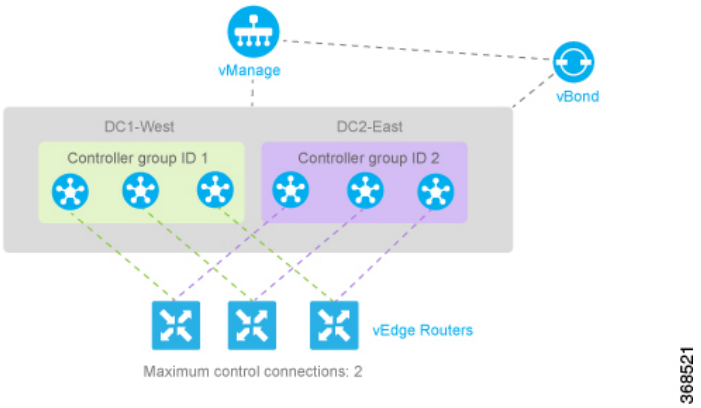

You configure the three vSmart controllers in DC1-West with controller group identifier 1:

```
vSmart-DC1(config)# system controller-group-id 1
```

The three vSmart controllers in DC2-East are in controller group 2:

```
vSmart-DC2(config)# system controller-group-id 2
```

We want all the Cisco vEdge device s to have a maximum of two OMP sessions, and we want each tunnel interface to have a maximum of two control connections and to not exclude any controller groups. So the only

configuration that needs to be done on the routers is to set the controller group list. We want Cisco vEdge device s in the west to prefer Cisco Catalyst SD-WAN Controllers in DC1-West over DC2-East:

```
vEdge-West(config)# system controller-group-list 1 2
```

Similarly, we want Cisco vEdge device s in the east to prefer DC2-East:

```
vEdge-East(config)# system controller-group-list 2 1
```

The software evaluates the controller group list in order, so with this configuration, the vEdge-West routers prefer vSmart controller group 1 (which is the West data center), and the vEdge-East routers prefer Cisco Catalyst SD-WAN Controller group 2.

You can fine-tune the controller group preference in other ways:

- Set the maximum number of OMP sessions allowed on the router to 1 (**system max-omp-sessions 1**). To illustrate how this works, let's look at a vEdge-West router. The router has only one tunnel interface, and that interface creates one control connection to Cisco Catalyst SD-WAN Controller list 1. If all the Cisco Catalyst SD-WAN Controllers in this group become unavailable, or if the connection between the router that the DC1-West data center goes down, the tunnel interface establishes one control connection to Cisco Catalyst SD-WAN Controller list 2, because this group is listed in the **system controller-group-list** command. If all Cisco Catalyst SD-WAN Controllers in both controller groups, or the connections to them, become unavailable, and if the vBond orchestrator also indicates that all these vSmart controllers are unreachable, the tunnel interface establishes a control connection to any other Cisco Catalyst SD-WAN Controller in the overlay network if other controllers are present.
- Set the maximum number of control connections that the tunnel interface can establish to 1 (**vpn 0 sdwan interface tunnel-interface max-control-connections 1**). Because the software evaluates the controller group list in order, for a vEdge-West router, this configuration forces the tunnel interface to establish a control connection to Cisco Catalyst SD-WAN Controller group 1. Again, if this controller group or data center becomes unreachable, the tunnel establishes a control connection with controller group 2, because this group is configured in the **system controller-group-list** command. And if neither controller group 1 or 2 is available, and if another vSmart controller is present in the network, the tunnel interface establishes a control connection with that controller.
- Exclude the non-preferred Cisco Catalyst SD-WAN Controller group for a particular tunnel. For example, for a vEdge-West router to prefer controller group 1, you configure **vpn 0 interface tunnel-interface exclude-controller-group-list 2** . As with the above configurations, if this controller group or data center becomes unreachable, the tunnel establishes a control connection with controller group 2, because this group is configured in the **system controller-group-list** command. And if neither controller group 1 or 2 is available, and if another Cisco Catalyst SD-WAN Controller is present in the network, the tunnel interface establishes a control connection with that controller.

Configure Redundant Control Connections on Single Device

When a Cisco IOS XE Catalyst SD-WAN device has two tunnel connections and the network has two (or more) data centers, you can configure redundant control connections from the Cisco vEdge device to Cisco SD-WAN Controllers in two of the data centers. It is recommended that do this using the minimum number of OMP sessions—in this case, two. To do this, you configure one of the tunnel interfaces to go only to one of the data centers and the other to go only to the second. This configuration provides Cisco SD-WAN Controller redundancy with the minimum number of OMP sessions.

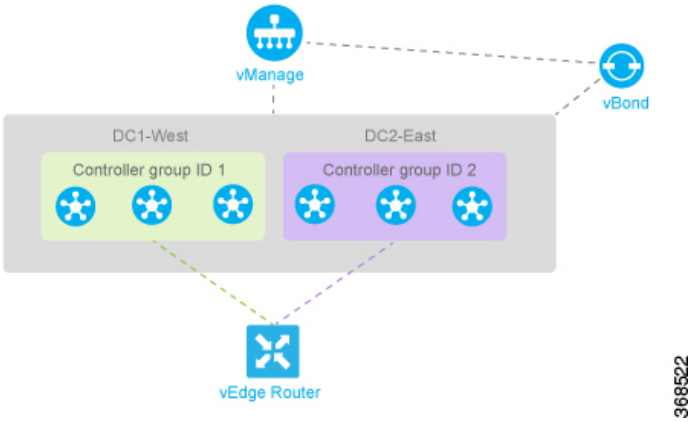

On the Cisco vEdge device router, define the controller group list and configure the maximum number of OMP sessions to be 2:

```
vEdge(config-system)# controller-group-list 1 2
vEdge(config-system)# max-omp-sessions 2
```

For one of the tunnels, you can use the default affinity configuration (that is, there is nothing to configure) to have this tunnel prefer a Cisco Catalyst SD-WAN Controller in group 1. You can also explicitly force this tunnel to prefer Cisco Catalyst SD-WAN Controller group 1:

```
vEdge(config-tunnel-interface-1)# max-control-connections 1
```

You do not need to configure **exclude-controller-group-list 2** , because the software evaluates the controller group list in order, starting with group 1. However, you could choose to explicitly exclude Cisco SD-WAN Controller group 2.

Then, on the second tunnel, configure it to prefer a Cisco SD-WAN Controller in group 2. As with the other tunnel, you limit the maximum number of control connections to 1. In addition, you have to exclude controller group 1 for this tunnel.

```
vEdge(config-tunnel-interface-2)# max-control-connections 1
vEdge(config-tunnel-interface-2)# exclude-controller-group-list 1
```

Configure Control Plane and Data Plane High Availability Parameters

This topic discusses the configurable high availability parameters for the control plane and the data plane.

Control Plane High Availability

A highly available Cisco Catalyst SD-WAN network contains two or more Cisco SD-WAN Controllers in each domain. A Cisco Catalyst SD-WAN domain can have up to 20 Cisco SD-WAN Controllers, and each Cisco vEdge device , by default, connects to two of them. You change this value on a per-tunnel basis:

```
vEdge(config-tunnel-interface)# max-control-connections number
```

Note The **no-control-connections** command is not applicable. Instead, use the **max-control-connections** command.

When the number of Cisco SD-WAN Controllers in a domain is greater than the maximum number of controllers that a domain's Cisco vEdge devices are allowed to connect to, the Cisco Catalyst SD-WAN software load-balances the connections among the available Cisco SD-WAN Controllers.

Note To maximize the efficiency of the load-balancing among Cisco SD-WAN Controllers, use sequential numbers when assigning system IP addresses to the Cisco vEdge devices in the domain. One example of a sequential numbering schemes is 172.1.1.1, 172.1.1.2, 172.1.1.3, and so forth. Another is 172.1.1.1, 172.1.2.1, 172.1.3.1, and so forth.

Data Plane High Availability

BFD, which detects link failures as part of the Cisco Catalyst SD-WAN high availability solution, is enabled by default on all Cisco devices. BFD runs automatically on all IPsec data tunnels between Cisco vEdge devices. It does not run on the control plane (DTLS or TLS) tunnels that Cisco SD-WAN Controllers establish with all Cisco devices in the network.

You can modify the BFD Hello packet interval and the number of missed Hello packets (the BFD interval multiplier) before BFD declares that a link has failed.

Change the BFD Hello Packet Interval

BFD sends Hello packets periodically to detect faults on the IPsec data tunnel between two Cisco vEdge devices. By default, BFD sends these packets every 1000 milliseconds (that is, once per second). To change this interval on one or more traffic flow, use the **hello-interval** command:

```
vEdge(config)#bfd color color hello-interval milliseconds
```

The interval can be a value from 100 to 300000 milliseconds (5 minutes).

Configure the interval for each tunnel connection, which is identified by a color. The color can be **3g**, **biz-internet**, **blue**, **bronze**, **custom1**, **custom2**, **custom3**, **default**, **gold**, **green**, **lte**, **metro-ethernet**, **mpls**, **private1**, **private2**, **public-internet**, **red**, or **silver**.

Change the BFD Packet Interval Multiplier

After BFD has not received a certain number of Hello packets on a link, it declares that the link has failed. This number of packets is a multiplier of the Hello packet interval time. By default, the multiplier is 7 for hardware routers and 20 for Cloud software routers. This means that if BFD has not received a Hello packet after 7 seconds, it considers that the link has failed and implements its redundancy plan.

To change the BFD packet interval multiplier, use the **multiplier** command:

```
vEdge(config)#bfd color color multiplier integer
```

Multiplier range: 1 to 60 (integer)

You configure the multiplier for each tunnel connection, which is represented by a color.

Control PMTU Discovery

On each transport connection (that is, for each TLOC, or color), the Cisco Catalyst SD-WAN BFD software performs path MTU (PMTU) discovery, which automatically negotiates the MTU size in an effort to minimize or eliminate packet fragmentation on the connection. BFD PMTU discovery is enabled by default, and it is recommended that you use BFD PMTU discovery and not disable it. To explicitly enable it:

```
vEdge(config)#bfd color color pmtu-discovery
```

With PMTU discovery enabled, the path MTU for the tunnel connection is checked periodically, about once per minute, and it is updated dynamically. With PMTU discovery enabled, 16 bytes might be required by PMTU discovery, so the effective tunnel MTU might be as low as 1452 bytes. From an encapsulation point of view, the default IP MTU for GRE is 1468 bytes, and for IPsec it is 1442 bytes because of the larger overhead. Enabling PMTU discovery adds to the overhead of the BFD packets that are sent between the Cisco vEdge devices, but does not add any overhead to normal data traffic.

If PMTU discovery is disabled, the expected tunnel MTU is 1472 bytes (tunnel MTU of 1500 bytes less 4 bytes for the GRE header, 20 bytes for the outer IP header, and 4 bytes for the MPLS header). However, the effective tunnel MTU might be 1468 bytes, because the software might sometimes erroneously add 4 bytes to the header.

Configure High Availability

CLI commands for configuring and monitoring high availability.

High Availability Configuration Commands

Use the following commands to configure high availability on a Cisco vEdge device:

```
bfd
  app-route
    multiplier number
    poll-interval milliseconds
  color color
    hello-interval milliseconds
    multiplier number
    pmtu-discovery
```

High Availability Monitoring Commands

show bfd sessions—Display information about the BFD sessions running on the local Cisco vEdge device.

Best Practices for Configuring Affinity

- In the **system controller-group-list** command on the Cisco vEdge device , list all the controller groups that are available in the overlay network. Doing so ensures that all the Cisco SD-WAN Controller in the overlay network are available for the affinity configuration, and provides additional redundancy if connectivity to the preferred group or groups is lost. You can manipulate the number of control connections and their priority with the maximum number of OMP sessions for the router, the maximum number of control connections for the tunnel, and the controller groups that the tunnel should not use (**exclude-controller-group-list** command).

  Listing all controller groups in the **system controller-group-list** command provides an additional layer of redundancy in situations where the Cisco vEdge device site is experiencing connectivity problems with the Cisco SD-WAN Controllers in the controller group list.

  To illustrate, consider a network with three controller groups (1, 2, and 3), and in which the controller group list on a Cisco vEdge device includes only groups 1 and 2 as preferred groups. In this scenario, if the router learns from the Cisco SD-WAN Validator that the Cisco SD-WAN Controllers in groups 1 and 2 are operational, but the router is unable to establish a connection to either device, it loses connectivity to the overlay network. However, if the controller group list contains all three controller groups and

group 3 is set up as a less preferred (excluded) group, the router still normally prefers groups 1 and 2, but would fall back and connect to the controllers in group 3 if it cannot connect to group 1 or group 2.

- The controller groups listed in the **exclude-controller-group-list** command must be a subset of the controller groups configured for the entire router, in the **system controller-group-list** command.
- When a data center has multiple Cisco SD-WAN Controllers that use the same controller group identifier, and when the overlay network has two or more data centers, it is recommended that the number of Cisco SD-WAN Controllers in each of the controller groups be the same. For example, if Data Center 1 has three Cisco SD-WAN Controllers, all with the same group identifier (let's say, 1), Data Center 2 should also have three Cisco SD-WAN Controllers, all with the same group identifier (let's say, 2), and any additional data centers should also have three Cisco SD-WAN Controllers.
- When a data center has Cisco SD-WAN Controllers in the same controller group, the hardware capabilities—specifically, the memory and CPU—on all the Cisco SD-WAN Controllers should be identical. More broadly, all the Cisco SD-WAN Controllers in the overlay network, whether in one data center or in many, should have the same hardware capabilities. Each Cisco SD-WAN Controller should have equal capacity and capability to handle a control connection from any of the Cisco vEdge device s in the network.
- When a router has two tunnel connections and the network has two (or more) data centers, it we recommend that you configure one of the tunnel interfaces to go to one of the data centers and the other to go to the second. This configuration provides Cisco SD-WAN Controller redundancy with the minimum number of OMP sessions.
- Whenever possible in your network design, you should leverage affinity configurations to create fault-isolation domains.