Improving Cisco ACI Virtual Edge Availability with VMware vSphere Proactive HA

Improving Cisco ACI Virtual Edge Availability

You can use VMware vSphere Proactive HA in vCenter 6.5 and later to improve Cisco ACI Virtual Edge availability.

Cisco Application Policy Infrastructure Controller (APIC) and VMware work together to detect a nonworking Cisco ACI Virtual Edge, isolate its host, and move its virtual machines (VMs) to a working host. Otherwise, if Cisco ACI Virtual Edge crashes, all its VMs can lose network connectivity.

You enable and configure vSphere Proactive HA in VMware vCenter, and in Cisco APIC, where the feature is called **Host Availability Assurance**. You can specify the amount of time that Cisco ACI Virtual Edge is not working before its host is quarantined and its VMs are moved.

Note

- Permission of the Cisco APIC account that you use for registration on VMware vCenter must have administrator rights or right to access the Cisco Application Centric Infrastructure (ACI) vCenter plug-in.
- vSphere Proactive HA is not available for Cisco ACI Virtual Edge when it is part of Cisco ACI Virtual Pod.
- For Host Availability Assurance to work, the VMware vCenter account used to create the Cisco APIC vCenter domain must have "Health Provider" write permission on the VMware vCenter.

How Improving Availability with vSphere Proactive HA Works

When you enable Host Availability Assurance, Cisco APIC creates a vSphere Proactive HA provider object in VMware vCenter. The object enables VMware vCenter to quarantine a host with a nonworking Cisco ACI Virtual Edge and move VMs out of that host. In Cisco APIC, you also specify how aggressively you want to trigger quarantine. You perform these tasks when you create a vCenter domain for Cisco ACI Virtual Edge.

When Host Availability Assurance is configured and enabled, Cisco APIC monitors Cisco ACI Virtual Edge on VMware vCenter. It uses the VMware vCenter inventory and OpFlex status to determine if Cisco ACI Virtual Edge is in a good or bad state. If Cisco APIC detects that Cisco ACI Virtual Edge is in a bad state, it tells VMware vCenter to put the affected host into quarantine.

VMware vCenter puts a host in quarantine mode according to one of three remediation modes, which you configure for the cluster:

- **Quarantine**: Hosts with health at yellow and red levels are put into quarantine mode.

  Note In Cisco ACI Virtual Edge Release 2.1(1), you can ensure that VM groups are moved out of Cisco ACI Virtual Edge hosts when the hosts stop working. The configuration overrides any affinity groups that would otherwise keep the VMs with particular hosts. For more information, see the section *VM Group Quarantine Protection*, in this guide.

- **Mixed**: Hosts with health at the yellow level are put into quarantine mode; hosts with health at the red level are put into maintenance mode.

  Note Although you can choose mixed remediation mode in VMware vCenter, the resulting behavior is the same as quarantine remediation mode.

- **Maintenance**: Hosts with health at yellow and red levels are put into maintenance mode.

  Important Do not choose maintenance mode remediation when you use vSphere Proactive HA. Maintenance mode requires that Cisco ACI Virtual Edge be powered off, which prevents the host from ever returning to a healthy state. Only use quarantine or mixed mode.

VMware vCenter also moves the VMs on that host to a host with a working Cisco ACI Virtual Edge. However, hosts in quarantine still might run data VMs if no healthy host is available, and any VM pinned by Distributed Resource Scheduler (DRS) rules to a quarantined host stays on the host. VMware vCenter also avoids moving any VMs to a quarantined host. However, you can deploy new VMs on a host in quarantine.

Benefits of Using vSphere Proactive HA

Using the vSphere Proactive HA feature is how you can detect and react to a Cisco Application Centric Infrastructure (ACI) Virtual Edge failure. If Cisco ACI Virtual Edge goes down, all of its virtual machines (VMs) that are connected to a VMware vCenter portgroup with **AVE** switching mode lose network connectivity.

vSphere Proactive HA also can prevent loss of connectivity in the following situations:

- **vSphere DRS**: The load-balancing vSphere Distributed Resource Scheduler (DRS) feature uses vSphere vMotion to automatically move VMs to enforce behavior that you define though affinity rules.

  However, DRS does not take Cisco ACI Virtual Edge into account. So in optimizing CPU and memory utilization, it can migrate a VM from a host with a working Cisco ACI Virtual Edge to a host where Cisco ACI Virtual Edge is not working.

- **Going into maintenance mode**: When you put a host into maintenance mode, DRS automatically migrates all of the host's VMs to another host. The host enters maintenance mode after all the VMs are moved.

  However, because Cisco ACI Virtual Edge is pinned to the host, DRS does not move Cisco ACI Virtual Edge, so the host doesn't enter maintenance mode. So without vSphere Proactive HA, you must power off the Cisco ACI Virtual Edge host for it to enter maintenance mode.

- **Going out of maintenance mode**: When you take a host out of maintenance mode, DRS can migrate VMs to that host because all of its CPU and memory are again available. However, Cisco ACI Virtual Edge must be powered on manually. This means that the Cisco ACI Virtual Edge might not be ready before DRS starts moving VMs back to the host.

  But vSphere Proactive HA lets Cisco ACI Virtual Edge power up on its own and delay moving VMs to the host until it is ready.

Important The host automatically enters and exits maintenance mode only in Cisco ACI Virtual Edge 2.1(1a) and later releases only. Earlier releases require that you put the host into and take it out of maintenance mode manually when using vSphere Proactive HA.

How vSphere Proactive HA Works

You enable and configure vSphere Proactive HA in VMware vCenter and in Cisco Application Policy Infrastructure Controller (APIC), where the feature is called **Host Availability Assurance**.

Enabling and configuring vSphere Proactive HA feature creates a vSphere Proactive HA provider object in VMware vCenter. The object enables VMware vCenter to quarantine a host with a nonworking Cisco ACI Virtual Edge and move VMs out of that host.

The feature also assigns a health status—green, yellow, or red—to every ESXi host in a Cisco ACI Virtual Edge virtual machine manager (VMM) domain. The status is green if the Cisco ACI Virtual Edge distributed virtual switch (DVS) is not added to the host or if the DVS is added and OpFlex is online. the status is yellow if the DVS is added and OpFlex is offline.

You also can specify how aggressively you want to trigger quarantine for the hosts.

Once you enable and configure vSphere Proactive HA, Cisco APIC and VMware vCenter work together to detect and isolate a nonworking Cisco ACI Virtual Edge.

1. Cisco APIC monitors Cisco ACI Virtual Edge on VMware vCenter.

   It uses the VMware vCenter inventory and OpFlex status to determine if Cisco ACI Virtual Edge is in a good or bad state. If Cisco APIC detects that Cisco ACI Virtual Edge is in a bad state, it tells VMware vCenter to put the affected host into quarantine by using the yellow level.

Improving Cisco ACI Virtual Edge Availability with VMware vSphere Proactive HA

2. VMware vCenter puts a host in quarantine mode according to a remediation mode, which you configure for the cluster in VMware vCenter:

Note Red status exists in VMware vCenter; it does not exist in Cisco APIC.

- **Quarantine**: Hosts with health at yellow and red levels are put into quarantine mode.

  Note In a Proactive HA cluster, VMware vCenter does not move a Cisco ACI Virtual Edge host into quarantine even if OpFlex goes down when an uplink or a physical network interface card (PNIC) is removed from the host.

- **Mixed**: Hosts with health at the yellow level are put into quarantine mode; hosts with health at the red level are put into maintenance mode.

  Note Although you can choose mixed remediation mode in VMware vCenter, the resulting behavior is the same as quarantine remediation mode.

  Note Do not choose maintenance mode remediation when you use vSphere Proactive HA. Maintenance mode requires that Cisco ACI Virtual Edge be powered off, preventing the host from ever returning to a healthy state. Use only quarantine or mixed remediation mode.

3. VMware Distributed Resource Scheduler (DRS) moves the VMs on the nonworking host to a host with a working Cisco ACI Virtual Edge.

Note Hosts in quarantine still might run data VMs if no healthy host is available, and any VM pinned by DRS rules to a quarantined host stays on the host. VMware vCenter also avoids moving any VMs to a quarantined host. However, you can deploy new VMs on a host in quarantine.

Note In Cisco ACI Virtual Edge Release 2.1(1) and later, you can ensure that VM groups are moved out of Cisco ACI Virtual Edge hosts when the hosts stop working. The configuration overrides any affinity groups that would otherwise keep the VMs with particular hosts. For more information, see and later the section VM Group Quarantine Protection, on page 7 in this guide.

4. Cisco APIC watches VMware vCenter events for hosts entering maintenance mode or rebooting and powers off Cisco ACI Virtual Edge when it is the only powered-on VM left on the host.
5. Cisco APIC powers on Cisco ACI Virtual Edge when the host is rebooted or taken out of maintenance mode.

Prerequisite for Configuring VMware vSphere Proactive HA

Complete the following task before you configure VMware vSphere Proactive HA.

Make sure that the VMware vCenter account used to create the Cisco APIC vCenter domain has "Health Provider" write permission on the VMware vCenter.

Enabling vSphere Proactive HA in Cisco APIC

Improving Cisco Application Centric Infrastructure (ACI) Virtual Edge in Cisco Application Policy Infrastructure Controller (APIC) consists of the following tasks:

- Enabling host availability assurance on the Cisco Application Centric Infrastructure (ACI) Virtual Edge VMM domain
- Specifying the time period before VMware vCenter quarantines any hosts on with Cisco ACI Virtual Edge has stopped working

You can perform these tasks in the Cisco APIC GUI when you create a vCenter domain for Cisco ACI Virtual Edge. See the section Create a VMM Domain Profile for Cisco ACI Virtual Edge in this guide for instructions.

You can perform these tasks with the NX-OS style CLI and REST API instead of the Cisco APIC GUI. See the sections Enabling vSphere Proactive HA Using NX-OS Style CLI and Enabling vSphere Proactive HA Using REST API in this guide.

Note When you add hosts to a cluster on which Proactive HA is already configured, and then add the host or attach the host to a Cisco ACI Virtual Edge VMM domain, those hosts may not work properly in some circumstances. The hosts may not work properly in Proactive HA or when Cisco ACI Virtual Edge or OpFlex goes down. The hosts also may not go into quarantine mode although the health status of the host is correctly set to yellow in Cisco Cisco Application Policy Infrastructure Controller (APIC).

To fix the problem, disable Proactive HA on the cluster and then re-enable it.

Enabling vSphere Proactive HA in VMware vCenter

Before you begin

Using vSphere Proactive High Availability (HA) requires VMware vCenter 6.5 or later.

Note When you add hosts to a cluster on which Proactive HA is already configured, and then add the host or attach the host to a Cisco Application Centric Infrastructure (ACI) Virtual Edge VMM domain, those hosts may not work properly in some circumstances. The hosts may not work properly in Proactive HA or when Cisco ACI Virtual Edge or OpFlex goes down. The hosts also may not go into quarantine mode although the health status of the host is correctly set to yellow in Cisco Cisco Application Policy Infrastructure Controller (APIC).

To fix the problem, disable Proactive HA on the cluster and then re-enable it.

Procedure

Step 1 Log in to the VMware vCenter Web Client.

Step 2 Choose **Home** > **Host and Cluster** > *cluster* > **Configure** > **Edit**.

Step 3 In the **Edit Cluster Settings** dialog box, choose **vSphere Availability** in the left navigation pane and then check the **Turn on Proactive HA** check box in the work pane.

Step 4 In the left navigation pane, choose **Proactive HA Failures and Responses**complete the following steps:

a) From the **Remediation** drop-down list, choose a remediation level.

Choose either **Quarantine**, which puts hosts with yellow and red levels into quarantine mode or **Mixed**, which puts yellow hosts into quarantine mode and red hosts into maintenance mode.

Note Do not choose **Maintenance**, which puts yellow and red hosts into maintenance mode. Maintenance mode requires that Cisco ACI Virtual Edge be powered off, which prevents the host from ever returning to a healthy state

b) Check the check box next to the vSphere Proactive HA provider to enable it.

The Cisco Application Policy Infrastructure Controller (APIC)-created provider would have "vmm-domain-name_APIC" as its name.

Manually Setting the Health Level of the ESXi Host

By default, the state of the VMware host is determined by the state of the Cisco Application Centric Infrastructure (ACI) Virtual Edge that resides on it.

You might want to override the default if you must do maintenance on the Cisco Application Centric Infrastructure (ACI) Virtual Edge. Setting the host state to yellow or red while Cisco ACI Virtual Edge is working properly puts the corresponding host into quarantine mode.

Or you might not want a specific host to go into quarantine, even if the Cisco ACI Virtual Edge on it goes down. Setting the state to green keeps the host active, disabling vSphere Proactive HA on the host.

Setting the health state manually to green prevents Cisco Application Policy Infrastructure Controller (APIC) from changing host status to yellow or red. You can view and set the health state using the Cisco APIC GUI, NX-OS style CLI, or REST API. See the section "Viewing and Setting a State on the Cisco ACI Virtual Edge Host Using the Cisco APIC GUI, on page 7." You also can view host health status and events in VMware vCenter. See "Tracking Health Updates for a Host in VMware vCenter, on page 7."

Viewing and Setting a State on the Cisco ACI Virtual Edge Host Using the Cisco APIC GUI

Before you begin

- You must have a host that contains Cisco ACI Virtual Edge.
- Host Availability Assurance must be enabled for the VMM domain on Cisco Application Policy Infrastructure Controller (APIC).

Procedure

Step 1 Log in to Cisco APIC.

Step 2 Go to **Virtual Networking** > **Inventory** > **VMM Domains** > **VMware** > ***VMM domain*** > **Controllers** and click the controller.

Step 3 In the **Controller Instance** work pane, in the **Health Policy** area, click the **+** (plus icon).

Step 4 Enter the host IP address and choose the state from the drop-down list, and then click **Update**.

Step 5 Click **Submit**.

Tracking Health Updates for a Host in VMware vCenter

If you enable Proactive HA, you can view events in VMware vCenter to track health updates for a host.

Procedure

Step 1 Log in to the VMware vCenter Web Client.

Step 2 Navigate to the host.

Step 3 In the central work pane, click the **Monitor** tab, **Tasks & Events**, and then **Events**.

The **Description** pane displays events for the host. In the **Type** column, VMware vCenter gives a warning for changes in the host's health, such as a degraded status or the host's subsequently entering a quarantine mode. There can be a 30-second delay between the report of a health problem and the host's being put into a different mode.

VM Group Quarantine Protection

When you enable Host Assurance Availability, the virtual machine (VM) reamains available even if Cisco Application Centric Infrastructure (ACI) Virtual Edge fails. Host Assurance Availability causes Cisco APIC to trigger the vMotion of VMs by setting the health status of the ESXi host where Cisco ACI Virtual Edge is in nonworking state to yellow or red.

However, Distributed Resource Scheduler (DRS) affinity rules and load balancing settings can cause VMs to stay or be placed on a nonworking host. Configuring protected VM groups enables Cisco APIC to autocreate anti-affinity rules in VMware vCenter, which force VMs part of the group to move out of nonworking host.

To protect VM groups, you must create the VM groups in VMware vCenter and enable protection for those VM groups in Cisco APIC. Each VM group should have all VMs that use Cisco ACI Virtual Edge.

You configure VM group protection in Cisco Application Policy Infrastructure Controller (APIC) on specific controllers. You can use the Cisco APIC GUI, NX-OS style CLI, or REST API.

Note For VM group quarantine protection to work, the VMware vCenter account used to create the Cisco APIC vCenter domain must have write privileges on the "Cluster" object in VMware vCenter.

Configuring VM Group Protection Using the Cisco APIC GUI

You can use the Cisco APIC GUI to configure VM group protection.

Before you begin

You must have configured VM groups in VMware vCenter and enabled vSphere Proactive HA in Cisco Application Policy Infrastructure Controller (APIC).

Procedure

Step 1 Log in to Cisco APIC.

Step 2 Go to **Virtual Networking** > **Inventory** > **VMM Domains** > **VMware** > ***domain*** > **Controllers** > ***controller***.

Step 3 In the controller work pane, choose the **Policy** and **General** tabs.

Step 4 In the **Protected VM Groups** area, check the check box for one or more VM groups.

Step 5 Click **Submit**.