



High Availability

The Cisco Catalyst SD-WAN Cloud OnRamp for Colocation solution allows various consumers to access various repetitive applications securely. The Cisco Catalyst SD-WAN Cloud OnRamp for Colocation solution High Availability (HA) is designed to handle several types of failure possible in a cluster deployment. The following types of failures can occur in a Cisco Catalyst SD-WAN Cloud OnRamp for Colocation solution deployment:

- Compute failure
- Switch failure
- Service chain failure

To resolve the failures, use the following mechanisms:

- Redundancy
- Failure detection
- [Redundancy, on page 1](#)
- [Handle Various Failure Scenarios, on page 5](#)

Redundancy

The following are the components where redundancy has been added to address failure of the component:

- x86 Compute Hardware—See [Redundancy of x86 Compute Hardware, on page 2](#).
- Network Fabric—See [Redundancy of Network Fabric, on page 2](#).
- Physical NIC/interface—See [Redundancy of Physical NIC or Interface, on page 2](#).
- NFVIS Virtualization Infrastructure—See [Redundancy of NFVIS, Virtualization Infrastructure, on page 2](#).
- Service-Chain/VNF—See [Redundancy of Service Chain or VNF, on page 2](#).
- Cisco Colo Manager—See [Recovery of Cisco Colo Manager, on page 5](#).

Redundancy of Network Fabric

Network Fabric—The hardware switch redundancy features are used to handle network fabric failures. In a switch failure, ensure that the standby switch takes over the traffic traversing through the failed switch.

Redundancy of x86 Compute Hardware

x86 Compute Hardware—Any hardware components such as, processor, storage, and others that are used on the x86 compute hardware can fail leading to a complete Cisco Cloud Services Platform (CSP) system failure. The Cisco vBond orchestrator continuously monitors the health of the x86 compute platform by using ICMP ping through the management interface. In a system failure, the orchestrator shows the device status and the service chains and VMs impacted. Take desired action to bring up service chains. See [Monitor Cisco Catalyst SD-WAN Cloud OnRamp for Colocation Solution Devices](#). Depending on the operational status of the VNFs (Virtual Network Function), the VMs must be brought up on a different CSP if enough resources are available. This action allows the VNF to retain the Day-N configuration. If the VNF disk is using local storage, the entire service group must be respun on another CSP device with the Day-0 configuration that is stored in the orchestrator.

Redundancy of Physical NIC or Interface

Physical NIC or interface—If a physical NIC (PNIC) or interface or cable fails or gets disconnected, the VNFs that are using these interfaces are impacted. If a VNF is using an OVS network, the port channel configuration is used to achieve a link redundancy. If a VNF is using an OVS network, and if the VNF has an HA instance, that instance has been already brought up on a different CSP. The failover happens to this VNF on the second CSP. If there is no second VNF instance, the service chain with the failed VNF must be deleted and reinstantiated.

Redundancy of NFVIS, Virtualization Infrastructure

Cisco NFVIS Virtualization Infrastructure—Multiple types of failures in the NFVIS software layer can occur. One of the critical components of CSP can crash or the host Linux kernel can panic or one of the critical components fails to respond. In case of critical component failures, the NFVIS software generates netconf notifications. The orchestrator uses these notifications to show the failure on vManage dashboard. If Cisco CSP or Cisco NFVIS crashes or control connection goes down, the orchestrator shows that device reachability is down. You can resolve a networking issue (if any), or reboot the CSP device. If device does not recover, you must proceed with removing the CSP device.

Redundancy of Service Chain or VNF

Table 1: Feature History

Feature Name	Release Information	Description
Placement of HA VNF NIC for Switch Redundancy	Cisco SD-WAN Release 20.5.1 Cisco vManage Release 20.5.1	This feature provides an optimum placement of service chains and therefore maximizes the resource utilization while accounting for switch redundancy. The VNICs of the HA primary and secondary instances are placed on alternate CSP interfaces to achieve redundancy at switch level.

Feature Name	Release Information	Description
Modifications to HA VNF NIC Placement	Cisco SD-WAN Release 20.6.1 Cisco vManage Release 20.6.1	This release modifies the placement of primary and secondary VNF VNICs on the physical NICs of the CSP device that are connected to redundant switch interfaces.

Service Chain or VNF—Some of the VNFs in the colocation service chain such as, firewall might support stateful redundancy features by using a standby VNF, whereas VNFs such as Cisco CSR1000V might not support stateful redundancy. The Cisco Catalyst SD-WAN Cloud OnRamp for Colocation solution relies on the VNFs to achieve VNF high availability. The HA support at service chain level isn't available. If a VNF supports stateful HA, it detects the failure and performs a switchover. The assumption is that the previously active VNF goes down and reboots as a standby VNF if the CSP device hosting the VNF is functional, and all the NIC or interface connectivity is functional. If the VNF isn't operational, the HA for VNF isn't functional from that time and you must fix the issue.


If a VNF doesn't support HA, it's assumed that the VNF reboots if any critical process fails within the VNF and no HA support is available for such VNFs.

Placement of Highly Available VNF NIC for Switch Redundancy

Starting from Cisco SD-WAN Cloud onRamp for Colocation Release 20.5.1, the network services in a service chain forward traffic without interruption even during switch failures. The traffic flow is uninterrupted because the virtual NICs (VNICs) of an HA virtual instance are placed on a different switch than the one that has the primary HA instance. For example, if VNF-primary is placed on the physical NIC of CSP1, which is connected to switch1, VNF-secondary is placed on the physical NIC of CSP2, which is connected to switch-2.

The image below shows the following:

- The solution provisions the primary instances of VNF #1 and VNF #2 to the SR-IOV ports on CSP #1, which are connected to switch #1.
- The secondary instances of VNF1 and VNF2 are placed on the SR-IOV ports of CSP2, which are connected to switch2.
- If switch #1 fails, the traffic continues to flow from the switch#2 of the first VNF and second VNF using the second switch.



Notes About HA VNF NIC for Switch Redundancy

- This feature applies to single-tenant clusters only, where the VNFs use SR-IOV interfaces, and where dual-homing to a switch is not supported. Multitenant clusters don't require this feature because they already use OVS interfaces, which are part of port channels and therefore, dual-homed to switches.
- The placement algorithm in the solution automatically places the service chains based on the redundancy requirements specified above. You don't need any manual configuration.
- When you upgrade Cisco vManage from earlier releases to Release 20.5.1, the following points apply when you use the HA VNF NIC redundancy feature:
 - For the new service groups that you create, the placement of the VNICs of an HA virtual instance on a CSP interface connecting to the alternate switch is automatic.
 - For existing service groups, detach the service group from a cluster, and then reattach it to the cluster to achieve switch redundancy for the service chain.
- At the time of placing the egress ports, the solution first attempts to place the egress port on the same CSP port that hosts the ingress VNF port. If the CSP port doesn't have sufficient bandwidth, the solution attempts to place the egress ports on the additional ports on the same CSP device that is connected to the same switch.

Starting from Cisco SD-WAN Cloud onRamp for Colocation Release 20.6.1, the solution supports service chains with bandwidth of up to 10 Gbps. The placement of the ingress and egress VNICs of a VNF could be on different CSP ports of the same CSP device if the bandwidth required is more than 5 Gbps and less than or equal to 10 Gbps.

Recommendations for Using Placement of HA VNF NIC for Switch Redundancy

- Design as many service chains as possible and provision these chains so that you use all service chain resources to the maximum capacity. This enables the colocation solution to utilize the VMs bandwidth completely in a sequential order without leaving any unused bandwidth on each port.
- Attach high-bandwidth service chains to colocation clusters, followed by the low-bandwidth service chains. For optimal resource utilization, attach highly available service chains to colocation clusters followed by the stand-alone service chains.

Recovery of Cisco Colo Manager

Cisco Colo Manager Recovery—Cisco Colo Manager is brought up on a CSP device in a Cloud OnRamp for Colocation. Cisco vManage selects a CSP with the DTLS tunnel to bring up Cisco Colo Manager. The Cisco Colo Manager recovery flow is required during the following scenario:

If a CSP hosting Cisco Colo Manager is considered for Return Material Authorization (RMA) process and there are at least two other CSP devices in the cluster after deleting this CSP, then a new Cisco Colo Manager is brought up automatically by Cisco vManage on one of the existing two CSP devices during a new configuration push.



Note You must power down the CSP device that has been considered for RMA process or perform a factory default reset on the CSP device. This task ensures that there is only one Cisco Colo Manager in the cluster.



Note A host with Cisco Colo Manager running can restart or reboot, and this action is not a recovery scenario as Cisco Colo Manager should come up intact with all the configuration and operational data.

If after a cluster is successfully activated and then Cisco Colo Manager becomes unhealthy, see [Troubleshoot Cisco Colo Manager Issues](#).

Handle Various Failure Scenarios

- VNF failure
 - If a VM in a service chain that is HA capable goes down, the standby VM takes over. This standby service chain is functional within few seconds. The Cisco NFVIS software on a CSP device tries to bring up the failed active VM if it's a monitored VM. If the VM recovers successfully, it switches to active and standby modes successfully. If the VM didn't recover successfully and you want to bring up HA capability on this VM, delete the service chain and bring up new service chain with HA capability. Here, VM detects that the failure is based on heartbeat and there must not be any impact on traffic (except few seconds). If an active VM recovers, this VM could become active again or stay as standby and this state varies from one VM to another.
 - If a VM isn't HA capable, the service chain fails and traffic is black holed. Cisco Colo Manager detects this failure and hence Cisco vManage as it receives notification that VM is down and service chain is down, Cisco vManage sends an alert. If the VM recovers successfully, the same notification

is sent and the service chain is functional without any intervention. If the VM doesn't recover successfully, delete the service chain and bring up a new service chain.

- Service chain failure

- If all VMs in a service chain support HA, service chains can have active and standby service chains. If an active service chain goes down, the standby service chain takes over and is functional within few seconds. This behavior is VM level HA and VM failover behavior takes over. Cisco NFVIS software on CSP also tries to bring up the failed active VMs (for monitored VMs) and if they recover successfully, the VMs switch over to active and standby modes successfully.
- If VMs aren't HA capable, the service chain fails and traffic is black holed. Cisco NFVIS and Cisco Colo Manager send notifications that VMs are down and Cisco vManage send an alert. Based on the notification, bring up another active service chain. If the service chain has recovered successfully, the same notification is sent and the service chain is functional without any intervention.

- Cisco CSP device failure

If a Cisco CSP is down, all the service chains and VMs running on that CSP are also down. Cisco Colo Manager sends notifications to Cisco vManage that the CSP device isn't reachable and Cisco vManage detects the DTLS connectivity loss with the CSP device. Cisco vManage sends alert about the CSP device and you must bring up the service chains on another CSP device by creating the service chains and pushing the configuration to a colocation. If there's not enough compute hardware, add another CSP device to a colocation and push the service chain configuration to the other CSP device.

Starting from Release 20.5.1, you can replace a faulty CSP device by creating a backup copy of the device in a colocation cluster. Therefore, when a CSP device fails, you can add a new CSP device to Cisco vManage, and restore the device to a state as the faulty device was in before the replacement. To know more about how to replace a CSP device, see [Return of Materials of Cisco CSP Devices](#).

- Switch link failure

If a link from a switch is down, the other switch takes over and service chain traffic continues.