



Cisco APIC-EM Multi-Host Support

- [Multi-Host Support, on page 1](#)

Multi-Host Support

A host is defined as an appliance, physical server, or virtual machine with Linux containers running instances of the Grapevine clients. The Grapevine root itself runs directly on the host's operating system and not in the Linux containers. You can set up either a single host or multi-host deployment. A multi-host deployment with three hosts is best practice for both high availability and scale. Each Grapevine root in a multi-host configuration maintains an Active/Active status with the other Grapevine roots and is therefore able to coordinate with the other Grapevine roots the overall management of the cluster.



Note Active/Active is defined as all Grapevine roots being operational and active.

Each host must be running the same controller software in the multi-host configuration. You are able to mix and match physical and virtual appliances in the multi-host configuration.

The multi-host configuration has the following requirements and features:

- Each host in a multi-host configuration requires a minimum of 32 GB of memory.
- A multi-host cluster comprised of 3 hosts is able to tolerate the loss of one of the hosts and supports a single fail-over (although with only two hosts, there is only software high availability, but no hardware high availability).



Note If a second host also fails in the three host cluster, the remaining host in the cluster will become inoperable and the cluster will go down. Therefore, in the event of the loss of one of the hosts, we recommend that you remove this host from the cluster using the configuration wizard and then either repair and rejoin this host to the cluster or join a new host to the cluster.

- As each host is configured with 32 GB of memory, if a host failure occurs then the remaining hosts would have a total 64 GB of memory which is sufficient to run the controller.
- All three hosts must reside in the same subnet.

Clustering and Database Replication

The clustering feature of the Cisco APIC-EM provides a mechanism for distributing processing and database replication among multiple hosts that run the exact same version of the controller. Clustering provides both a sharing of resources and features, and enables system high availability and scalability.

Security Replication

In a multi-host environment, the security features of a single host are replicated among the other two hosts, including any X.509 certificates or trustpools. Once you join a host to another host or to a cluster, the Cisco APIC-EM credentials are shared and become the same as that of the host you are joining or the pre-existing cluster. The Cisco APIC-EM credentials are cluster-wide (across hosts) and not per-host.

**Note**

We strongly suggest that any multi-host cluster that you set up be located within a secure network environment. For this release, privacy is not enabled for all of the communications between the hosts.

Service Redundancy

The Cisco APIC-EM provides high availability support using service redundancy. A Cisco APIC-EM cluster can be set up across multiple Linux containers within multiple hosts. On each host, the Grapevine root is an application running on the host and the Grapevine clients are created and reside in the containers. Both the Cisco APIC-EM services and database are then instantiated across the clients within the Linux containers:

- Cisco APIC-EM Services:
 - For service high availability, if a service fails then Grapevine (the Elastics Service Platform) spins up a new instance to replace it. If Grapevine is unable to spin up the new instance on the same container after a sole instance fails, then it spins up a new container and then spins up the new instance on this container.
 - Cisco APIC-EM supports a replacement service instance model. For example, assume that one of the roots on a single host spins up an instance. If that host and its root goes down, then another host on another root spins up an instance to ensure continuity of that service.
- Cisco APIC-EM Database:
 - The Cisco APIC-EM services use a PostgreSQL database management system. PostgreSQL has a built-in master-slave model for synchronizing data across replicated databases to respond to any failover situation.
 - The master and slave postgres instances are grown across different Linux containers and across different hosts. The data of these postgres instances are synchronized using PostgreSQL's built-in data streaming replication mechanism. With three hosts, there is one master (with a master postgres instances) and two slaves (each with a slave postgres instance).
 - If the master fails, then the slave seamlessly takes over.
 - In the event of a failure by the master, an election process occurs among the remaining hosts to determine which becomes the new master. This election process can also be triggered by resetting the controller using the CLI or rebooting the host.

**Caution**

To protect against any hardware failure, you must deploy the Cisco APIC-EM on a cluster with three hosts.

Multi-Host Synchronization

Whenever there is a configuration change on one of the hosts, Grapevine synchronizes the change with the other two hosts. The supported types of synchronization include:

- Database—Synchronization includes any database updates related to the configuration, performance, and monitoring data.
- File—Synchronization includes any changes to the configuration files.

Multi-Host Monitor Process

Grapevine is the main component that manages high availability operations in a cluster. To ensure proper cluster high availability operation, Grapevine uses both health checks and heart beats.

Health checks are used to monitor processes that are low performing and not running properly. Services that run on Grapevine have health checks that are periodically invoked. If there is any indication of an unhealthy service, Grapevine will harvest and regrow that service.

In addition to the health checks, Grapevine also uses heart beats between the services, clients, and roots to monitor the status of the cluster. Grapevine monitors these heart beats for any processes that may have failed. If there is no heart beat, then this indicates that a process has failed and to correct for this situation, Grapevine regrows the service.

Grapevine also uses a heart beat to monitor for adequate memory and storage capability for the cluster. If a heart beat indicates that the cluster's memory or storage fails below an appropriate level necessary for successful operations, then Grapevine will not grow any new services.

Split Brain and Network Partition

When Cisco APIC-EM is configured as a multi-host cluster, a private network connection is set up between the hosts. This private network connection is used by each host to monitor the health and status of the other cluster hosts. A split brain occurs when there is a temporary failure of the network connection between the hosts, for example, due to any of the following occurrences:

- Physical disconnection of the network connection from a host
- Loss of power to one or more hosts
- Cisco APIC-EM appliance failure

During a split brain occurrence, situations can arise where each separate host is sending commands to a given network device without any coordination with the other hosts, and the results can be problematic.

To correct for a split brain event, when the private network connection fails between one of the hosts, the other two hosts create a quorum and establish a network partition between themselves and the failed host with the following results:

- The split brain or network partition scenarios are handled by ensuring quorum (majority reads and rights) to the controller database.
- The side of the partition with the "minority" stops operating, since it is unable to perform quorum (majority reads and rights) to the controller database.
- The side of the partition with the "majority" continues to operate, since they are *able* to perform quorum (majority reads and rights) to the controller database.