



Clustering Avi Controllers That Are in Different Networks

Avi Technical Reference (v16.3)

Clustering Avi Controllers That Are in Different Networks

[view online](#)

Avi Controller clusters provide high availability (HA) and redundancy, as well as increased scale for analytic workloads. A common question is how to best architect the placement of Avi Controllers to ensure the highest redundancy.

Avi Controllers communicate with each other over their single management IP address, the [cluster IP address](#). They also use this path to communicate with all Avi Service Engines (SE) within the fabric.

Avi Controllers are not required to exist within the same IP network, but there are a few limitations worth noting:

- Avi Controllers must be within the same region (ideally the same data center). This is because they need to quickly synchronize databases and perform actions such as log indexing and data retrieval. For example, an Avi Vantage user who logs onto Avi Controller A may be attempting to access logs that are first indexed on Avi Controller C, and may experience some latency.
- Avi Controllers have the option of sharing a [cluster IP address](#). The cluster IP address is owned by the primary Avi Controller within the cluster. In order to share an IP address, all Avi Controllers must have a NIC in the same network.
- Each Avi Controller must have access to the IP addresses of the other Avi Controllers through configured network routes.

Note: OpenStack requires Avi Vantage to maintain a cluster IP address. For this reason, Avi Vantage deployed into an OpenStack cloud does not support clustering of Avi Controllers from different networks.