



# Avi GSLB Architecture

Avi Technical Reference (v20.1)

Copyright © 2020

# Avi GSLB Architecture

[view online](#)

This article has been replaced with [Avi GSLB Architecture, Terminology, and Object Model](#).

Content prior to 29Aug2018:

Avi GSLB provides simplified and centralized configuration and monitoring of global applications. In a typical environment, the corporate name server delegates one or more subdomains to Avi GSLB, which then owns these domains, and provides responses to DNS queries from clients. Avi GSLB provides an active/backup model for backup or disaster recovery applications, and an active/active model to respond with the most optimal site, based on load, proximity, etc. This article overviews the architecture, functionality, and terminology for Avi GSLB.

## Features

- [DNS-based load balancing](#)
- Active/DR GSLB
- Active/active GSLB with consistent hash and [geolocation](#)
- [Support for NAT-aware \(public-private\) IPs](#)
- Hybrid cloud support
- [Centralized provisioning](#)
- [Control-plane and data-plane health monitoring](#)
- Multitenancy support
- Support for 3rd-party load balancers
- [GSLB site selection \(17.1.5\)](#)
- Support for a two-level GSLB algorithm: location-based at the global-service level with other algorithms at the pool level (17.2.3)
- [Support for GSLB site persistence based on HTTP cookies \(17.2.5\)](#)
- [For GSLB site selection ...](#)
- Ability define a single fallback site and set the preferred-site option (17.2.5)
- Up to 16 fallback sites (17.2.7)

## GSLB-related articles:

- [Avi DNS Architecture](#)
- [Avi GSLB Overview](#)
- [Avi GSLB Service and Health Monitors](#)
- [Avi GSLB Site Configuration and Operations](#)
- [NAT-aware Public-Private GSLB Configuration](#)
- [GSLB Site Cookie Persistence](#) is available with 17.2.5.
- [Selective Assignment of a GSLB Service to DNS Virtual Services](#)
- [GSLB Site Selection with Fallback and Preferred-Site Options](#)
- [GSLB Integration with F5 GTM](#)
- [Upgrades in an Avi GSLB Environment](#)

## Architecture and Terminology

### Global Applications

A global application (represented as a [GSLB service](#) in Avi Vantage) is an application which cannot or should not be wholly contained by one Avi Vantage deployment, three of the most common reasons being

- There exists a geographically dispersed community of users (clients),
- Resilience to loss of a data center is required, and
- The occasional migration to or addition of another data center.

Global server loading balancing (GSLB) is the act of balancing a global application's load across individual Avi Vantage (and possibly third-party ADC) deployments. In this document, we refer to these deployments as [GSLB sites](#).

### Functionality

A global application requires a solution capable of performing these functions:

1. Definition and ongoing synchronization/maintenance of the GSLB configuration,
2. Monitoring the health of configuration components,
3. DNS-based steering of inbound application requests to the best source for responses, and
4. The actual processing of those responses.

In Avi's design, an individual GSLB site need not perform all four functions. For example, a virtual service defined by Controller\_cluster\_A (site\_A) might respond to requests steered to it by DNS functionality overseen by Controller\_cluster\_B (site\_B).

To better understand the steps one takes and the objects one touches to implement the Avi GSLB solution, it is important to keep in mind these distinct GSLB functions.

### GSLB Site Types

GSLB sites fall into two broad categories:

- An Avi site runs an Avi Vantage Controller cluster and SEs.
- An external site runs third-party ADCs from vendors such as F5, Citrix, etc.

With the four functions in mind, we can further categorize an Avi [GSLB] site as either an active or a passive GSLB site.

- Active sites perform some combination of the above-listed functions #1-#3, typically all three.
- Passive sites only perform task #4, i.e., the hosting of virtual services that respond to requests from the clients of global apps.
- A passive site is never asked to furnish steering information (function #3), so none is ever pushed to it.
- A passive site's health is determined by a health monitor running on an active site. It is unaware of and has no means to determine the health of other sites.
- An active site may host virtual services that back-end global apps. To that extent, it is additionally serving in a *passive* role; we do not say it is more "active."

### Leader and follower *active* sites

- Exactly one active site is statically designated as the GSLB leader. It is the one on which the initial GSLB site configuration is performed.

- The other active sites are GSLB followers. What GSLB configuration data they need is propagated to them from the leader.
- The only way to switch leadership is through an override configuration from a follower site. This override can be invoked in the case of site failures or maintenance.
- Centralized analytics are rolled up by and available from the GSLB leader. Localized metrics and logs are available for the DNS virtual services hosting the GSLB records.

### Avi sites versus external sites

- Avi sites can perform any combination of the four functions. Each Avi site is characterized/identified by an Avi Controller cluster.
- External sites can only perform function #4. The opacity of an external site prevents it from participating in a GSLB site configuration, which only contains information about *Avi Controller clusters*, such as their address and credentials.

Note: Not every virtual service homed to a particular Avi or external site need participate in the Avi GSLB solution. Put another way, not every application is a global application.

### Defining a GSLB Service

A GSLB service is the representation of a global application. The corresponding [GslbService object](#) identifies

- the application service's name,
- the FQDN of the application,
- one or more GSLB pools comprising virtual service members running the application in the GSLB sites,
- the priority/weights of the GSLB pools,
- the weights of the virtual service members *within* those pools, and
- the monitoring methods to be used to make sure members are alive.

The above merits further explanation.

A GSLB (pool) member is typically a virtual service (as opposed to a service running on a solitary back-end server without an ADC front-ending it). Like other virtual services, a GSLB member is represented by an IP address:port or name.

A GSLB pool is a collection of GSLB pool members sharing the same priority, but potentially different weights.

Note: A GSLB pool is much different than an Avi server pool. The former aggregates *virtual services*, the latter aggregates *servers*. That said, there is a hierarchical relationship between the two types of pools.

GSLB service supports minimum members which refers to the minimum number of servers within GSLB service pools. \* For GSLB service with *priority* as the GSLB service algorithm, the pools are sorted based on the priority and the minimum servers' logic is applied. \* For GSLB service with *geo-location* as the GSLB service algorithm, the pools are sorted based on the client's location and the minimum servers' logic is applied.

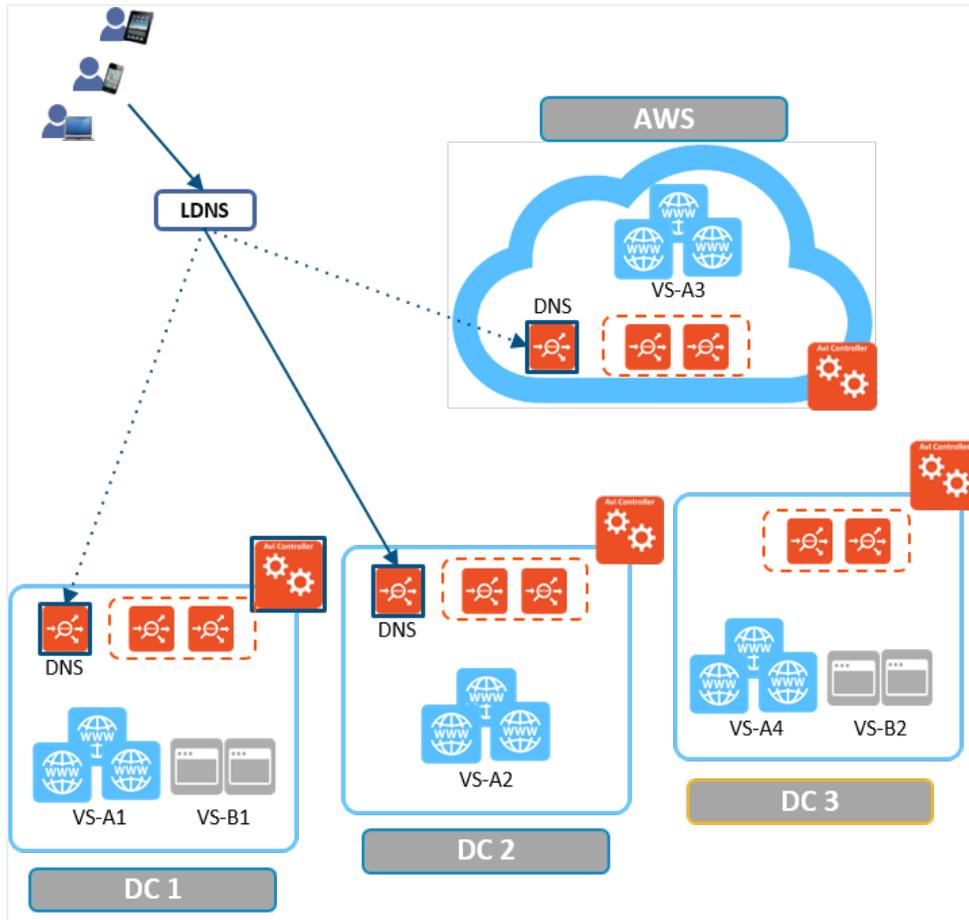
For instance, in the case of two pools ? P1 and P2 with priority ten and five respectively, the traffic will be sent to P1 as long as it is up. With minimum servers logic, P1 will be chosen exclusively as long as the number of servers active in P1 matches the defined minimum server value. That is, if the minimum servers' value is set to five, then P1 will be chosen as far as it has five active(up) servers.

The [GslbHealthMonitor](#) object is the GSLB counterpart of the [virtual service] HealthMonitor object. It provides configuration for monitoring the health of GSLB pool members.

An active/backup configuration is used for backup/disaster recovery. One GSLB pool could be configured as the active pool and another GSLB pool as standby. If/when all virtual services in the active pool go down, the GSLB DNS instead directs global-app traffic to the (former) standby pool.

There exist a variety of active/active configurations. All are used to put more resources to work as well as increase application availability.

- A load-unaware active/active configuration incorporates multiple virtual services across sites to be equally responsible for handling the application. Two options are offered ? (weighted) round robin and consistent hash. Consistent hash is the one option that can provide persistence.
- A load-aware active/active configuration steers clients to the most optimal site based on the observed load of the members. This option is slated for a future release.
- A geo-location-aware active/active configuration steers clients to the closest site. Member selection is based on the proximity of the client to the members. Refer to [this article](#).



In the above example, Santa Clara, Chicago, Austin and NY are Active sites while Boston is a passive site. Among the active sites, Santa Clara performs the function of GSLB leader, redirecting DNS queries to the appropriate followers (Chicago, Austin and NY).

## GSLB Services

A GSLB service is the representation of a global application, that front-ends applications across multiple sites. This configuration object hosts the FQDN of the application, the backing virtual services in various sites, and the priority/ratios for the various members. This object also hosts the monitoring methods to make sure that the group members are alive.

For more information, please refer to [Avi GSLB Service and Health Monitors](#).

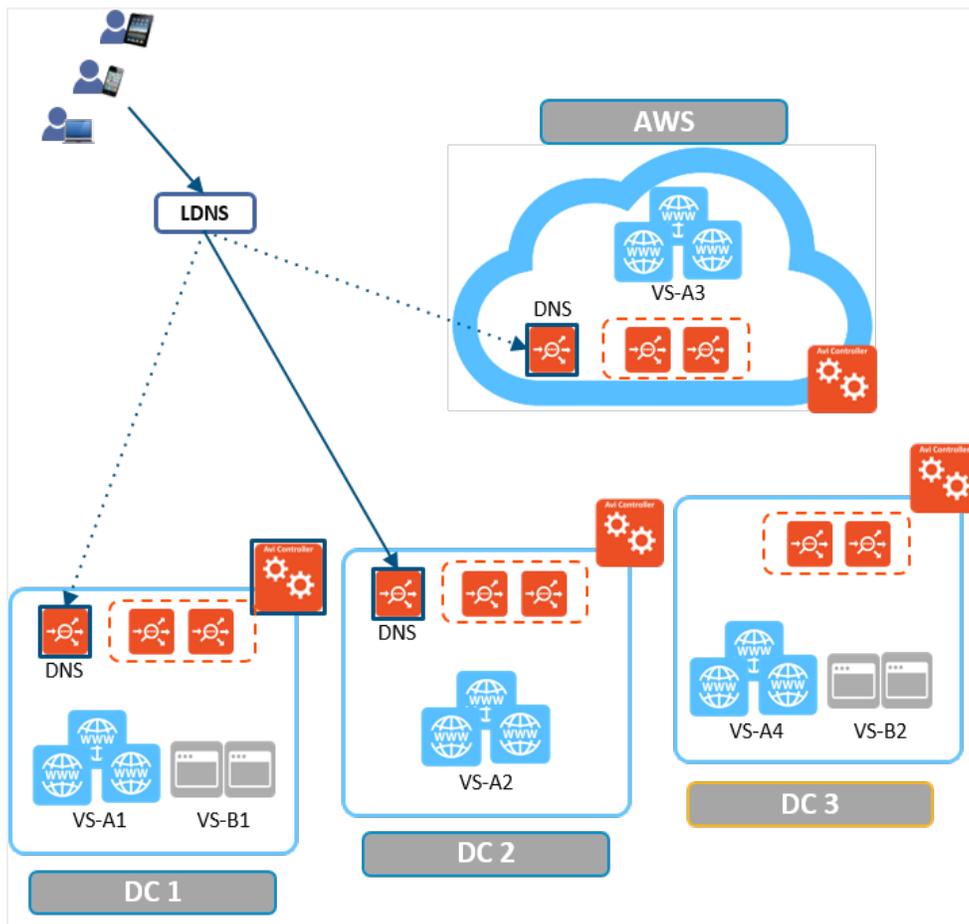
## Health Checks

The three options for GSLB member health checks are:

- Control-plane health - The DNS virtual service collects the results of health checks performed by remote Controllers on their local member virtual services.
- Data-plane health - The DNS virtual service performs GSLB health checks to member virtual services.
- Both control- and data-plane health - For a GS member to be marked as UP, both control- and data-plane health should report UP. If control-plane health check is failing (due to some remote Controller being inaccessible) but data-plane health checks work, the GS member is marked as UP.

For more information, please refer to [Avi GSLB Service and Health Monitors](#).

## Overall Flow



A GSLB service is the representation of a global application that front-ends applications across multiple sites. This configuration object hosts the FQDN of the application, the backing virtual services in various sites, and the priority/ratios for the various members. This object also hosts the monitoring methods to make sure that the group members are alive.

## Site Failure Handling

In case of Site failure, one of the following case occurs:

- GSLB leader site fails
- Follower site fails

In both cases, the DNS service marks all GS members of the failed site as DOWN. In the above example, NY site has crashed while Santa Clara (GSLB leader) and Chicago sites are still functioning.

In case of GSLB follower site failure, the GSLB configuration is synchronized with currently UP site. In case only the Controller at a follower site fails, then the site's DNS server will not receive GSLB configuration updates.

In case of a GSLB leader site failure, no modification can be made to the GSLB configuration.

## Site Failure Recovery

All active sites detect successful connection to a follower through the control-plane health monitor updates. In the case of control-plane health monitors, active sites query all the other sites to fetch health and load information of all virtual services which are behind GSLB services.

The DNS service then updates the GS member status of the recovered sites as UP or DOWN, based on the GS health monitors.

- Follower Site Recovery

GSLB leader detects successful connectivity to follower and updates it with the updated GSLB configuration.

- Leader Site Recovery

Only an admin can resume updates to the GSLB configuration.

To learn more about changing the GSLB leader and site recovery, please refer to [Avi GSLB Site Configuration and Operations](#).

## Tenancy and GSLB

1. The `Avi admin` user can do a GSLB site configuration.
2. The `Avi admin` user or a tenant administrator can configure the GSLB DNS virtual service.
3. Each tenant user can create a GSLB service and health monitors for it. The requisite DNS records are registered on the GSLB DNS virtual service configured in step 2.

### Notes

- A tenant administrator has visibility only to his/her tenant's GSLB services.
- The tenant should have read access to the `Gslb` configuration object.