



# Avi GSLB Architecture and Terminologies

Avi Technical Reference (v20.1)

Copyright © 2020

# Avi GSLB Architecture and Terminologies

[view online](#)

## Overview

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, terminology and object model for Avi GSLB.

## Global Application

A global application requires a solution capable of performing the following functions: \* Defining, synchronizing, and maintaining GSLB configuration. \* Monitoring the health of the configuration components. \* Optimizing application service for clients by providing GSLB DNS responses to their FQDN requests. \* Processing global application requests.

## Key Functions

A GSLB solution should be capable of performing the following four key functions:

1. Definition and ongoing synchronization/maintenance of the GSLB configuration ? an Avi Controller responsibility
2. Monitoring the health of configuration components ? a responsibility shared by Avi Controllers and Service Engines
3. Optimizing application service for clients by providing GSLB DNS responses to their FQDN requests based on the GSLB algorithm configured ? the responsibility of an Avi GSLB DNS running in one or more Service Engines
4. Processing of application requests ? the responsibility of services placed on Avi SEs and/or running on third-party servers/load balancers.

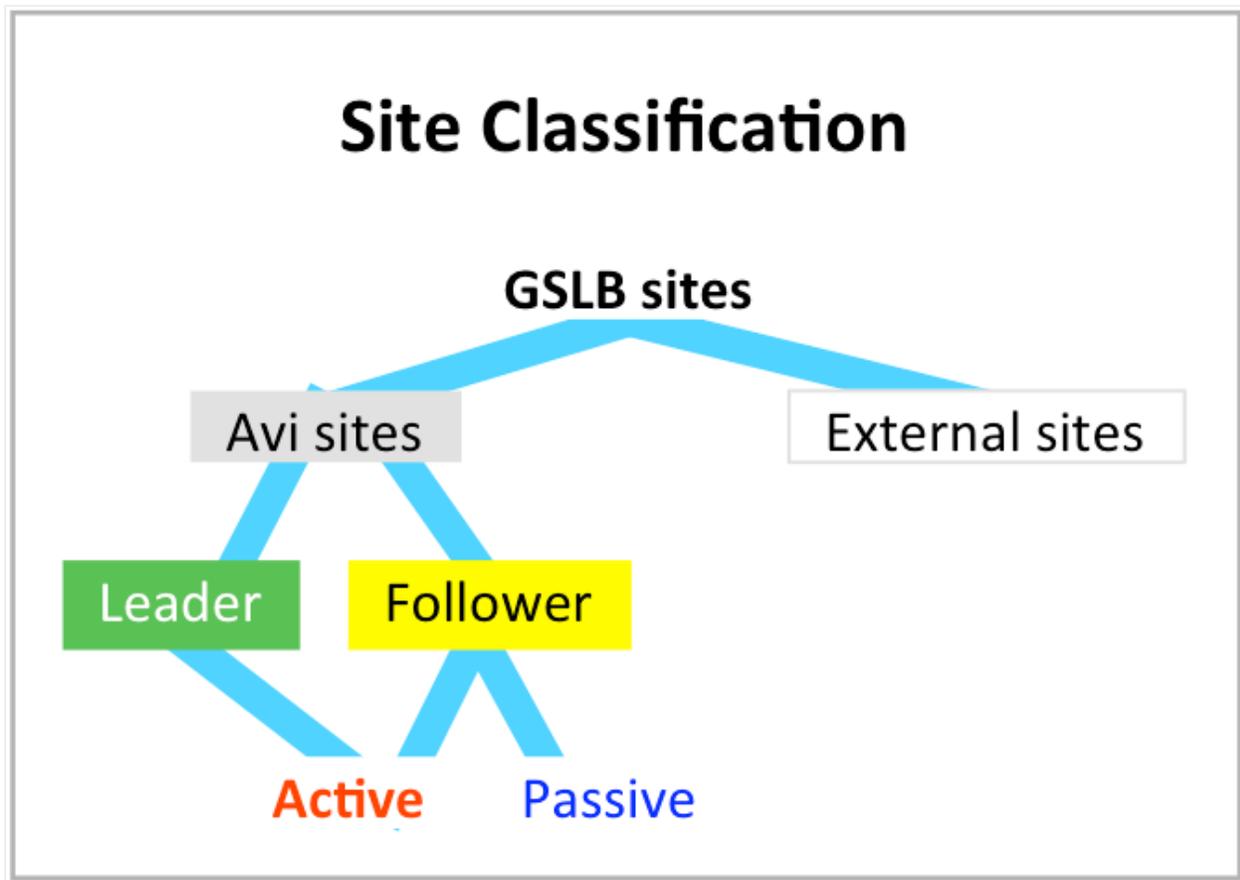
In Avi GSLB, an individual GSLB site does not perform all four functions. For example, a virtual service defined by one of the sites might respond to requests steered to it by GSLB DNS functionality overseen by Controller by the other site.

## Key Entities

Key entities involved in Avi GSLB are : \* GSLB Sites \* Avi DNS \* GSLB Service \* GSLB Pool \* GSLB Pool Members \* Load Balancing Algorithm \* Health Monitors

## GSLB Sites

GSLB sites are of the following two types: \* Avi Sites: Leader and follower (active and passive) sites \* External Site



**External site**

An external site runs third-party ADCs. Such sites can only perform the function 4 as mentioned above. 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 an Avi site or external site need to participate in the Avi GSLB solution. This means not every application is a global application.

**Avi Sites**

An Avi site has an Avi Vantage Controller cluster and SEs and can perform any combination of the four key functions.

**Active and Passive Sites**

The Avi GSLB site is further classified as either an active or a passive GSLB site.

- Active site ? Active sites perform some combination of the above-listed functions 1-3, typically all three.
- Passive site ? 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 requested to furnish steering information (function no. 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.

### Leader and Follower Active Sites

- The Avi site on which the administrator first defines a GSLB configuration is automatically designated as the GSLB leader. Exactly one active site is statically designated as the GSLB leader.
- The other active sites subsequently added are GSLB followers. The GSLB configuration 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. Read this article for more information.

### Leader Active Site

The leader active site is responsible for the functions 1, 2, 3, and 4 as mentioned in the previous section.

Followers are further classified as active or passive, based on their behavior vis-a-vis the first three functions.

Note: Starting with Avi Vantage 18.2.5, GSLB leader periodically attempts to resynchronize the objects which are in error state. The default value for the resynchronization interval is 300 seconds.

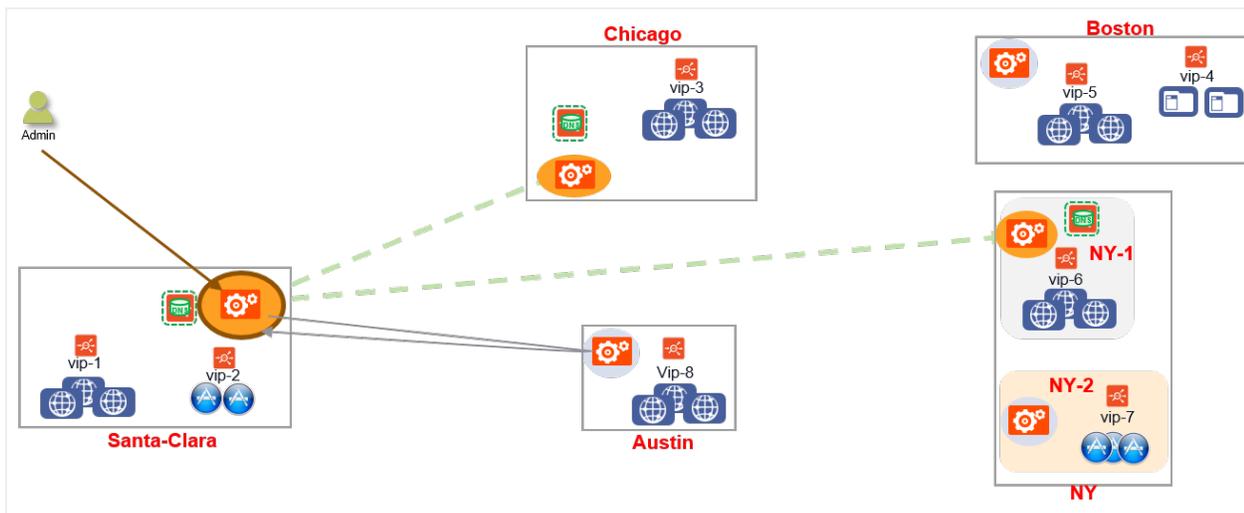
### An Active Follower

- Receives the GSLB configuration from the leader, and thus can take over leadership in the event of the leader's failure.
- Must actively monitor the health of other GSLB sites.
- May host an authoritative DNS for the Avi global applications defined by Avi GSLB. Such redundancy makes DNS service for the GSLB configuration more reliable and boosts the performance of address resolution.

### A Passive Follower

- Does not receive the GSLB configuration, and thus can not take over for a failed leader.
- Does not monitor other sites. Its health is determined by a health monitor running on an active site.
- Does not have DNS participating in the GSLB configuration. It may run DNS for applications unrelated to the GSLB deployment.

Note: An Avi site may participate in exactly one Avi GSLB configuration. If *site\_A* is a participant in *GSLB\_config\_1*, an attempt to incorporate *site\_A* into *GSLB\_config\_2* results in an error.



In the above example, Santa Clara, Chicago, and NY-1 are active sites. They each have a DNS service running. Boston, Austin, and NY-2 lack DNS services; they are passive sites.

## Avi DNS VS

Avi Vantage DNS runs a virtual service with the application profile type of *System-DNS* and a network profile using per-packet load balancing. Avi DNS VS will be authoritative for the subdomains delegated to Avi for GSLB.

For more information on Avi DNS , refer to [Avi DNS Architecture](#).

## GSLB Service

A GSLB service is the representation of a global application. It front ends instantiations of the application that are deployed at multiple sites. Key elements of a global service include:

- The global application service's name, by which administrators reference the GSLB configuration.
- The FQDN of the application, by which end-user clients reference it. A list of domain names can be provided for aliasing (e.g., `www.foo.com` and `foo.com`).
- Backing services running in various GSLB sites, organized into one or more GSLB pools. Most of the time, a backing service will be an Avi virtual service. However, it can be a third-party ADC's virtual service, or even an IP:port on some solitary server.  
Note: A *GSLB* pool is much different from a *server* pool. The former aggregates *backing services*, while the latter aggregates *servers*.
- Health monitoring methods by which unhealthy components can be identified so that alternatives may be selected.
- A time-to-live (TTL) ranging from 1-86400 seconds that determines the frequency with which clients need to obtain fresh steering information for client requests. If none is specified for a particular GSLB service, the TTL value defaults to the one specified in the DNS application profile (default there is 30 seconds).

Note: It is recommended to use caution when using very low values of TTL since some DNS or operating systems will discard a very low TTL.

## GSLB Pools

A GSLB service can comprise of multiple GSLB pools. A GSLB service can switch load across multiple GSLB pools based on the priority of the pools as explained in the example below or based on the geolocation.

Note: A GSLB pool is different from a server pool. The former aggregates back-end services, while the latter aggregates servers.

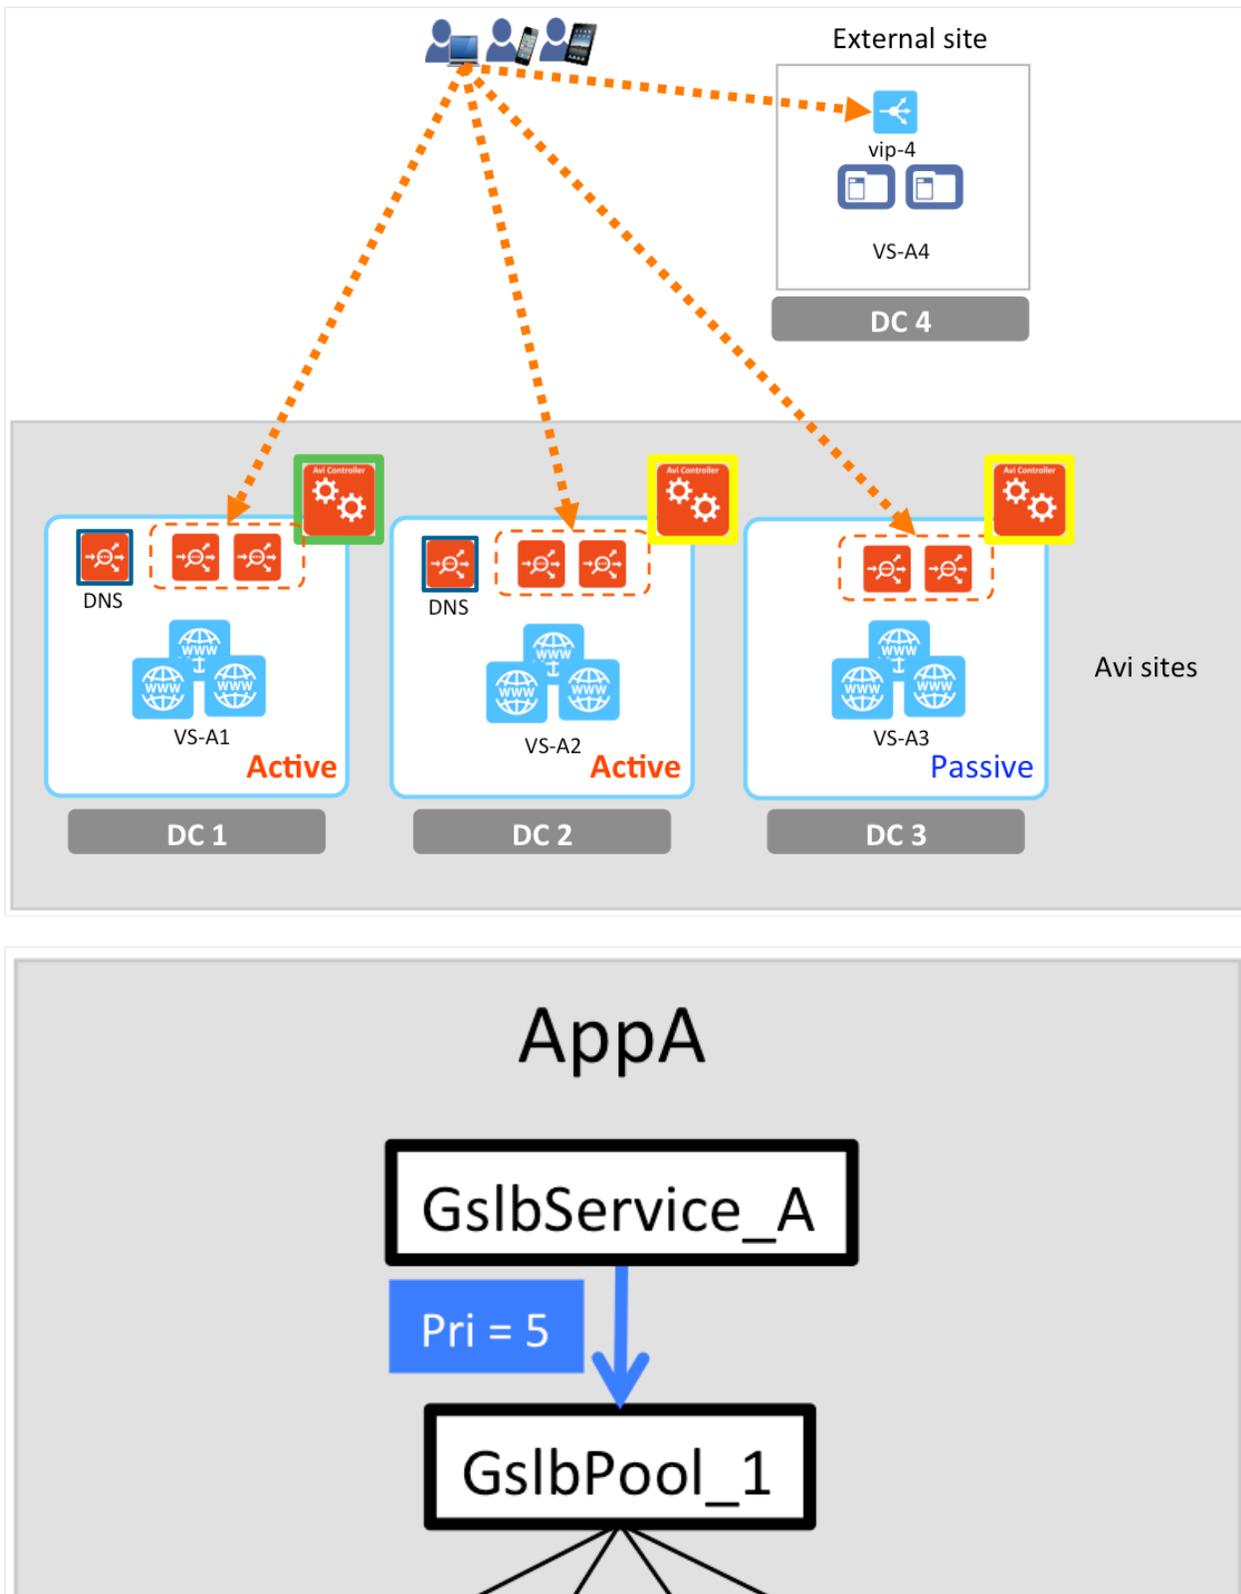
- `GSLB_SERVICE_ALGORITHM_GEO` ? A GSLB pool is selected based on the geolocation. When a GSLB service is optioned for `GSLB_SERVICE_ALGORITHM_GEO`, pool priority holds no significance.

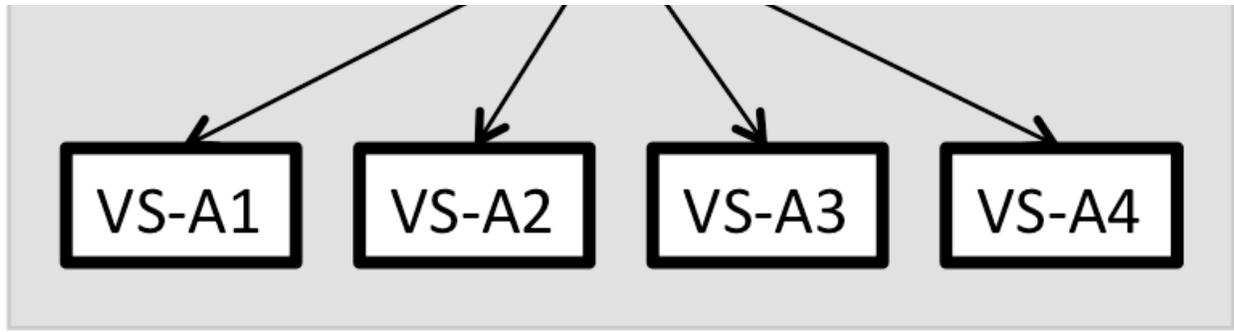
For more information, refer to [Geolocation Based Load Balancing Algorithm](#).

- `PRIORITY` ? A GSLB pool is selected based on its pool priority.

Refer to the below diagram. For *AppB*, which corresponds to the *GslbService\_B* is comprised of two pools, with priorities 5 and 10. As long as the pool of highest priority (*GslbPool\_3*) is up and not at its connection limit, all traffic will be directed to it. *GslbPool\_2* will remain idle. However, if a pool is inaccessible, down, or at maximum capacity, a lower priority pool will be chosen instead.

A global application, *AppA*, spanning four sites, *DC1* through *DC4*, is depicted below. The GSLB pool object *GslbPool\_1* aggregates services *VS-A1* through *VS-A4* across which load will be balanced.





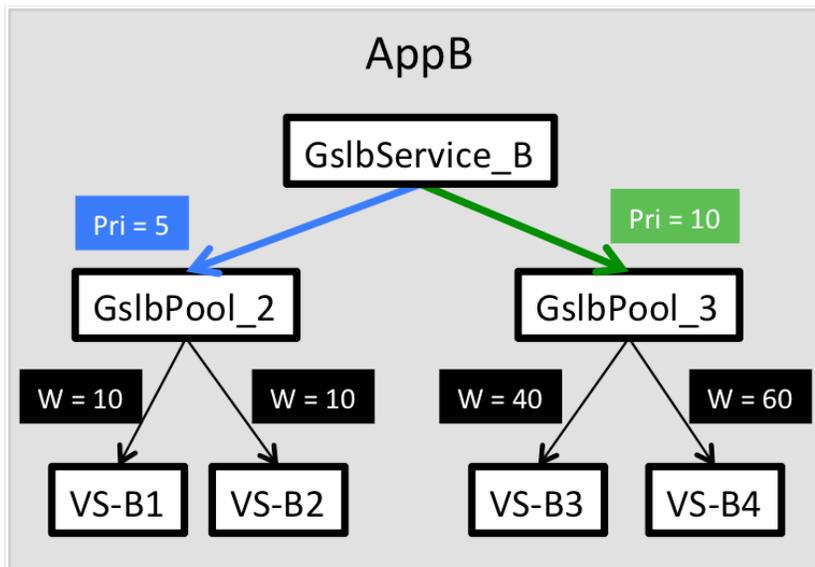
### GSLB Pool Members

The services that comprise a GSLB pool (e.g., *VS-B3* and *VS-B4*) are called GSLB pool members. Members can be specified by the followings:

- The virtual service name
- An IP address, to specify standalone servers or VIPs defined by third-party load balancers, and/or
- A DNS name, for example, to specify a DNS-based load balancer such as AWS ELB.

### GSLB Pool Member Weights

All members of a GSLB pool share the same priority, but each member of the pool may potentially have different weights. When a pool is selected and its load is distributed across the members via the round-robin algorithm, these weights dictate the fraction directed to each member. In the example depicted below, as long as *VS-B3* and *VS-B4* are healthy and able to accept the load, *GslbPool\_3*'s higher priority will direct all load to it, but weights will cause Avi Vantage to direct 40 percent of that load to *VS-B3* and 60 percent to *VS-B4*.



### Load Balancing Algorithms for GSLB Pool Members

Once a particular pool has been selected, a GSLB algorithm (as indicated in the `GslbPool.algorithm` parameter) balances load across the pool's member services.

- Round robin? Traffic is distributed evenly across all members and the traffic can optionally be skewed by weight. These are implied by the value of the member's `GslbPoolMember.ratio` parameter, which defaults to 1, and may range from 1 to 20. For example, if virtual services A, B, and C have ratios of 1, 2 and 3 respectively, virtual service A will receive one-sixth, B will get one-third, and C will get one-half the load.
- Consistent hash ? Load is distributed based on the client's source-IP address (likely a DNS resolver address). If EDNS processing is enabled, the source IP address will be found in the [ECS option](#). An integer mask ranging from 1 to 31 can be applied to the client IP address, if there are multiple local DNS in a given network, in one site. This algorithm can provide persistence. [GSLB site cookie persistence](#) is the other.
- Geolocation-based ? Client requests are directed to the optimal site based on the longitude and latitude of the client relative to the GSLB members. Refer to [Geolocation-based Load Balancing Algorithm for GSLB Members](#) for details.
- Topology-based algorithm ? For more information, refer to [Topology-Based GSLB Algorithm](#).

[Starting with Avi Vantage release 18.2.5, the support to select n records from m GSLB pools has been provided. Refer to the Selecting n Number of Pool Members from m Number of Pools](#) for more details.

## Avi GSLB Health Monitors

Using health-monitoring option available for GSLB, unhealthy components (GSLB service) can be identified and the request is forwarded to remaining pool members.

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

## Additional Information

- Bypassing the Load Balancing Algorithms

Under certain circumstances, it may be desirable to bypass the load balancing algorithm for certain clients based on certain conditions being met. Refer to the [GSLB Site Selection with Fallback and Preferred-Site Options](#) article.

- Minimum members to be UP

The GSLB service minimum members parameter (introduced in 17.2.4) modifies the first step in the selection process by taking into account the number of *UP* member services within a given GSLB service pool. If a pool's count of *UP* services drops below the minimum members, traffic is not directed to the pool, even though the pool might otherwise have been selected (based on either its priority or proximity).

- Two-level GSLB Pool Member Service Selection  
Avi CLI and Avi UI configuration of this feature is documented in [this how-to article](#).
- Prior to 17.2.13, each GSLB pool within a GSLB service was assigned a distinct priority. Multiple GSLB pools with a GSLB service may be assigned the same priority. If a set of pools shares a priority that exceeds the priority of the other pools, then Avi Vantage selects among that set in a round-robin fashion, as long as the minimum members logic is satisfied. The valid priority range was expanded to include a priority of zero. Setting a pool's priority to zero effectively disqualifies it from the selection. This might be done temporarily, for example, to perform maintenance on the services comprising the pool.

**Note:** Setting pool priority to 0 is not equivalent to disabling it. A zero-priority pool will continue to be health monitored; its state will be displayed as *UP* or *DOWN*, not *DISABLED*.

- **Firewall Settings**

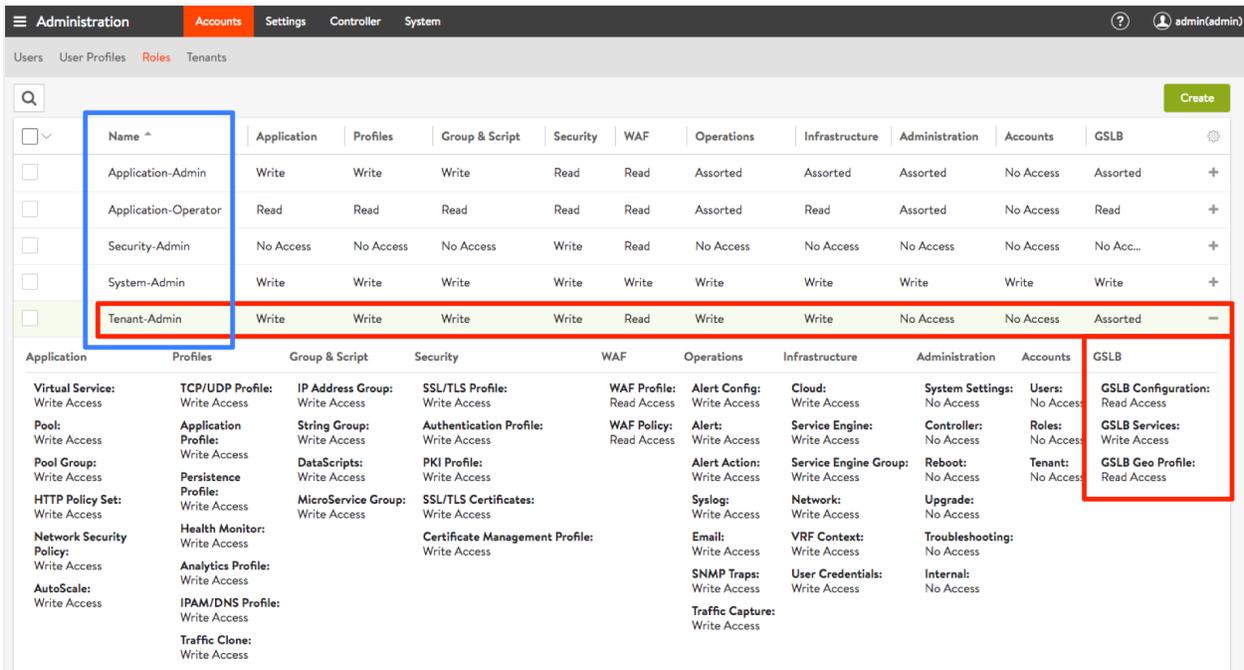
Firewall settings must permit bi-directional Controller communication among all active sites. In addition, for data-plane health monitors to work, all active sites need to be able to probe all virtual services at every site participating in a GSLB service and configured to be probed (a GSLB service can be configured to probe all members or just non-Avi members).

## Tenancy

The expected isolation and administrative restrictions of a multi-tenant architecture extend to Avi GSLB.

## Predefined Roles

Navigate to Administration > Accounts > Roles to check the default predefined roles, These roles are assigned to authorized Controller users. In our particular example, clicking the plus sign at the right end of the Tenant-Admin role row expands all ten columns, revealing more details about different aspects of the Avi Vantage platform. GSLB rights (red rectangle) focus on access to three GSLB entities ? the GSLB configuration, global applications (services), and the geolocation database that supports geolocation.



For each role in the system, access rights can be independently set. The below table summarizes GSLB access rights for three of the predefined roles. Additional, custom-defined roles can be imagined.

Predefined Roles	Access Rights	
<b>System Admin</b>	GSLB configuration	Write access to entire GSLB configuration
	GSLB services	Write access to all global apps in all tenant
	GSLB geolocation database	Write access to geolocation database
<b>Tenant-Admin</b>	GSLB configuration	Read access to the GSLB configuration relevant to the tenant
	GSLB services	Write access to all global apps in all tenant

	GSLB geolocation database	Read access to geolocation database
<b>Application-Admin</b>	GSLB configuration	Read access to entire GSLB configuration
	GSLB services	Write access to all global apps in all tenant to which this used is assigned
	GSLB geolocation database	Read access to geolocation database

### GSLB Administrators, Global Application Administrators

An authorized user may enjoy a different set of access rights, depending upon the requirements.

GSLB may only be configured or reconfigured by a user account having write access to GSLB Configuration. Such access is included in the definition of the predefined `System-Admin` role, the role assigned to the user named `admin`. GSLB configuration can only be performed from within the tenant named `admin`. A system administrator places the DNS virtual service in the `admin` tenant or some other tenant. However, a site's GSLB DNS virtual service can only be configured into a *non-admin-tenant* via the Avi CLI using the UUID of the DNS virtual service. Each application administrator associated with the Avi tenant can then use the shared-DNS virtual service for their global application(s).

A user with write access to GSLB Services and read access to GSLB Configuration can define global applications (GSLB services) in their own tenant. The requisite DNS records are registered on the shared-DNS service. The tenant administrator is able to get analytics of the GSLB services defined only for that tenant.

Note: Starting with Avi Vantage release 18.2.6, GSLB service creation are no longer dependent on the tenant user is logged in, but instead on the permissions of the GSLB user.

### High Availability Recommendations for GSLB

For high availability, it is recommended to configure DNS for GSLB on a Service Engine group that is scalable to two or more Service Engines. It is also recommended to implement DNS for GSLB in more than one location.

This can be implemented in the following two ways:

1. Have at least two geographically separated active GSLB sites. For each site configure DNS onto a scalable SE group.
2. If only one active site is defined, ensure a minimum of at least one geographically remote cloud. On that remote cloud, configure DNS for GSLB on a scalable SE group. Also define all virtual services to support the mission-critical applications running on the origin location.