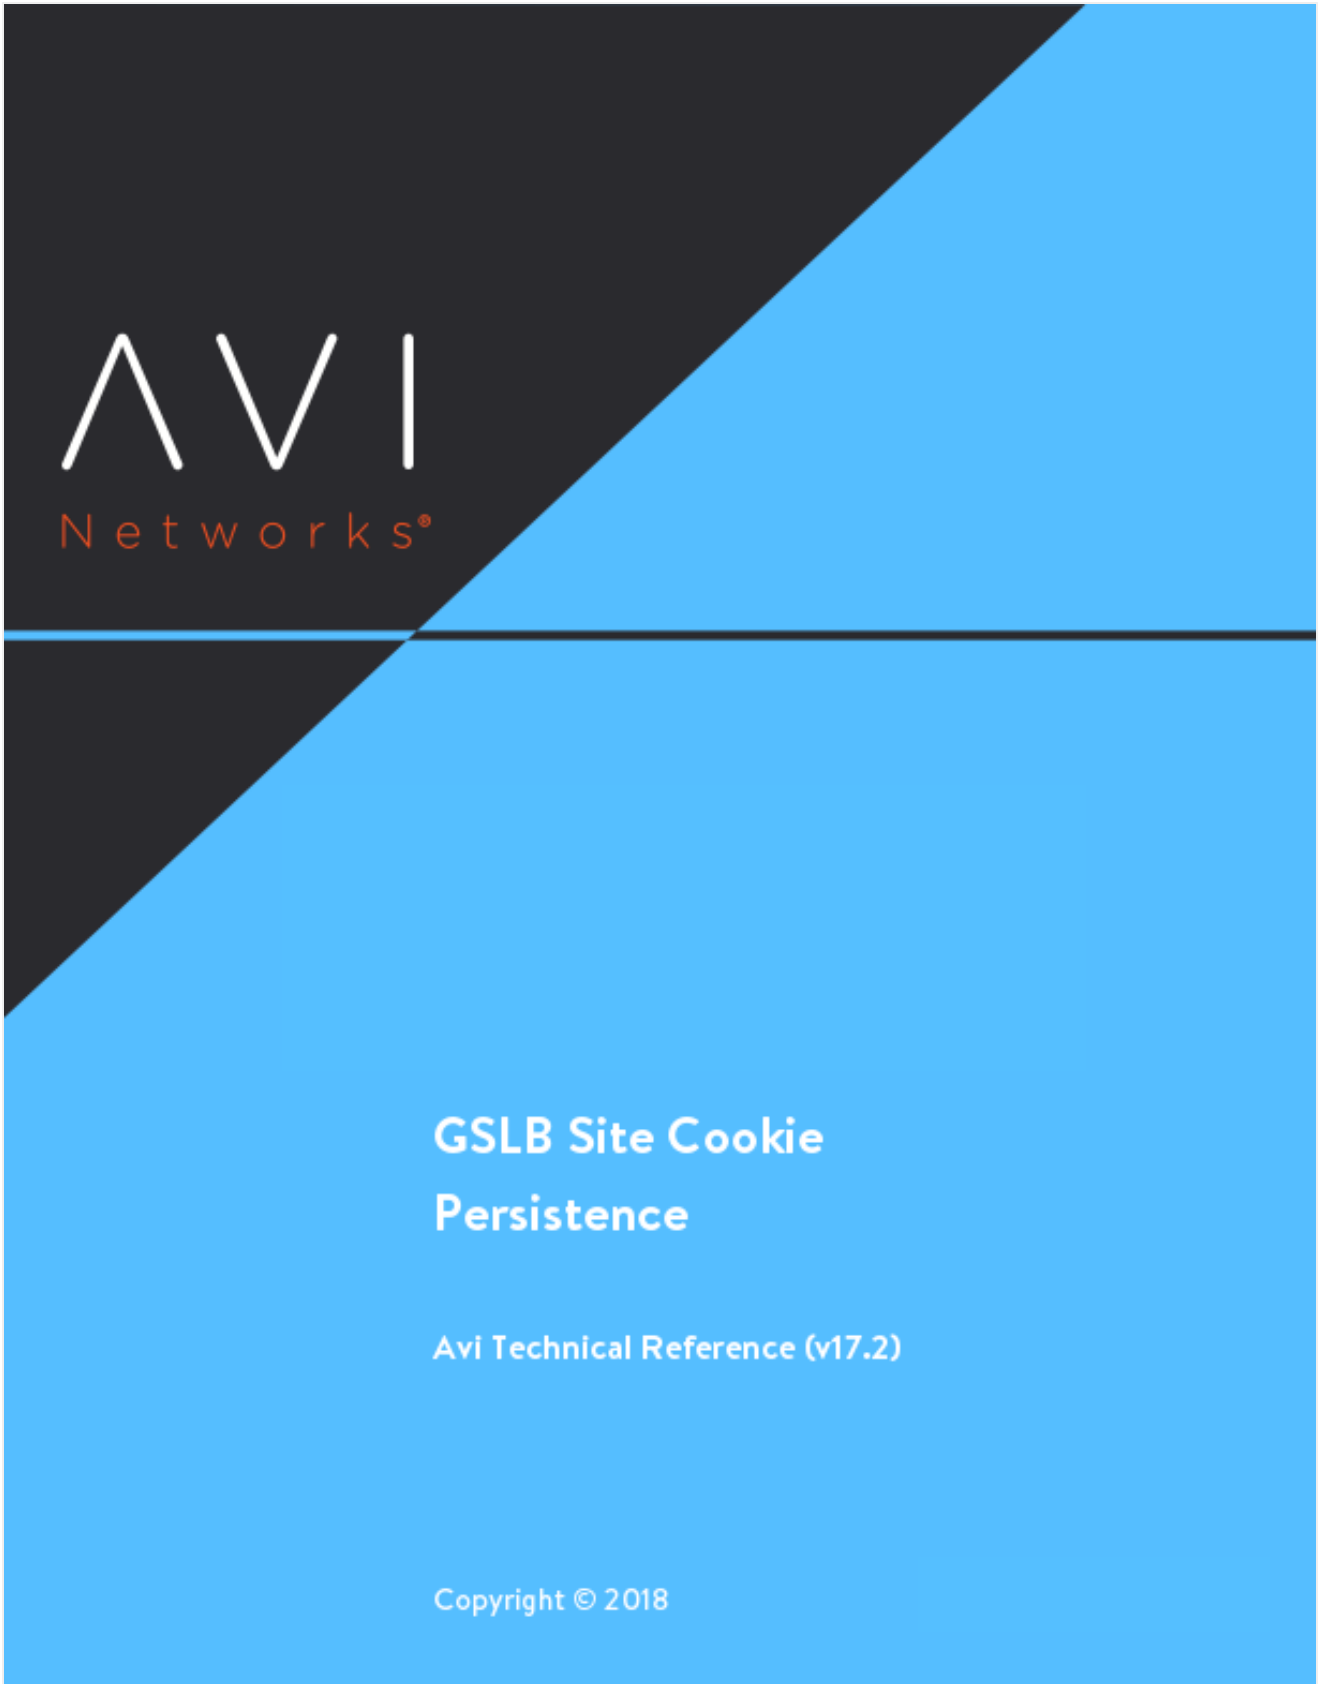

GSLB Site Cookie Persistence

Avi Technical Reference (v17.2)

GSLB Site Cookie Persistence

view online

Starting with Avi Vantage release 17.2.5, long-lived transactions from clients in a GSLB application can be configured to persist to the sites in which their transactions were initiated. This feature is implemented using HTTP site cookies created by Avi Service Engines.

Overview

Some applications require stickiness between a client and a server. That is to say, all requests in a long-lived transaction from a client must be sent to the same server; otherwise, the application session may be broken, with negative impact on the client. This is accomplished by turning on GSLB site cookie persistence, which takes precedence over the configured GSLB algorithm.

In an active-active GSLB deployment, site persistence is extremely important. Typically, site persistence is not an issue in active-standby deployments.

Restrictions

Avi Vantage checks for the below-listed conditions and will emit appropriate error messages if violations are attempted. The need for these restrictions is better understood after reading this entire article.

- Site persistence applies only to Avi VIPs; non-Avi (aka third party) VIPs cannot participate.
- Site persistence across multiple virtual services within the same Controller cluster is not supported.
- For site persistence to be turned on for a global application, all of its individual members must run on active sites. Conversely, a site cannot transition from active to passive if an Avi GS member participating in a site-persistent GSLB service runs on it.
- For site persistence to work, an Avi GS member must be unique across all GSLB services. That is to say, it cannot be a GSLB pool member in more than one GSLB service.
- A site-persistence pool is an internal pool construct created by the Controller and associated with GSLB VS members when site persistence is turned on. Users may not perform or change this association. Avi's pool group feature can't be configured for site-persistence pools.
- The `is_federated` option is added to the required PKI profile to ensure the profile can be replicated across all GSLB members. There can be *only one* `is_federated` PKI profile defined. Because there is only one, there is no need to explicitly associate the federated PKI profile with any GSLB service.
- Federated PKI and application persistence profiles cannot be
- Associated with unfederated profiles.
- Created if GSLB is not turned on.

How Cookie-Based Site Persistence Works

Refer to figure 1, which depicts Phase 1 in the life of a long-lived transaction.

1. The client asks its DNS resolver to resolve `x.foo.com`.
2. The corporate DNS determines that there are two authoritative DNSs (at Site1 and Site2), and recommends that the DNS resolver try the Site1 DNS.
3. The DNS at Site1 receives the DNS resolver's query and ? using whatever global load balancing algorithm is in force ? recommends VS1 at Site1 to the DNS resolver and sets the TTL.
4. In turn, the DNS resolver passes the recommendation and TTL on to the client.

5. Each SE in the group implementing VS1 is aware that site persistence is ON. The client's first request contains no site cookie, so the SE creates one and passes it back. The cookie is AES-256-encrypted and based on the `cluster_uuid` and `vs_uuid` strings. Below is an example of such a cookie:

```
Set-Cookie: FOO=1S509ceebd-0913-4aomuTiRccdU0ujbfY6eCVkL9muOBwIsnT5fhrMTMMM4-fapeQ2SEGb3ny69-iJQYG6Xg6SmLq9x7cr
```

6. The two-way dialog indicated by the double-ended blue arrow continues as long as SEs in the group see this cookie.

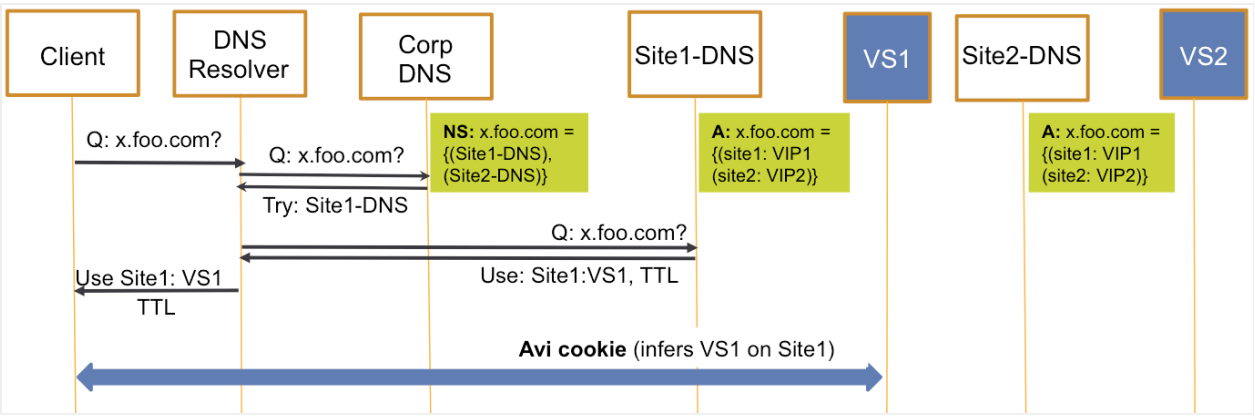

Figure 1. Phase 1 of GSLB cookie-based site persistence

Figure 2 shows what happens some time, later, after the TTL has expired.

1. That expiration forces the client to once again ask the DNS resolver to provide an IP address for `x.foo.com`.
2. The DNS resolver once again requests the corporate DNS to provide an authoritative DNS. From its cache, the corporate DNS happens to recommend Site2's DNS.
3. The client queries the Site2 DNS and this time provides VIP2, the IP of address of VS2, which is local to it.

In figure 3 we see the client initiating a dialog with VS2 at Site2. Unbeknownst to it, the cookie previously obtained accompanies that request.

1. An SE at Site2 receives the request with site cookie attached. It decrypts the cookie, and immediately can tell this request is part of an ongoing conversation that did not start on its site. Rather, the conversation needs to be proxied to *VS1* at *Site1*.
2. In proxying the request to VS1, VS2 passes the request to it, making sure to set the return address to itself.
3. The SE responds to the client, using content provided by VS1 at Site1.

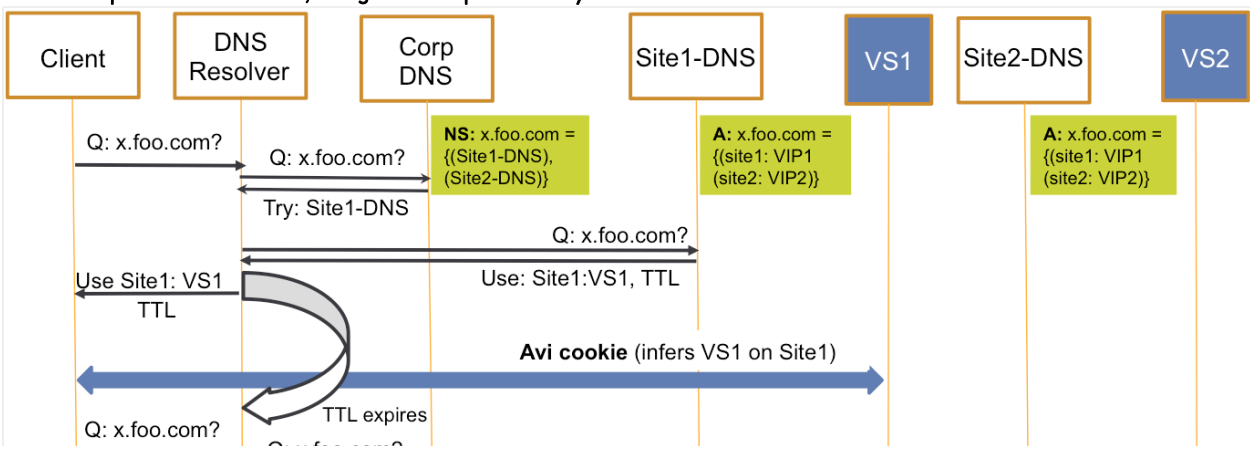

Figure 2. Phase 2 of GSLB cookie-based site persistence

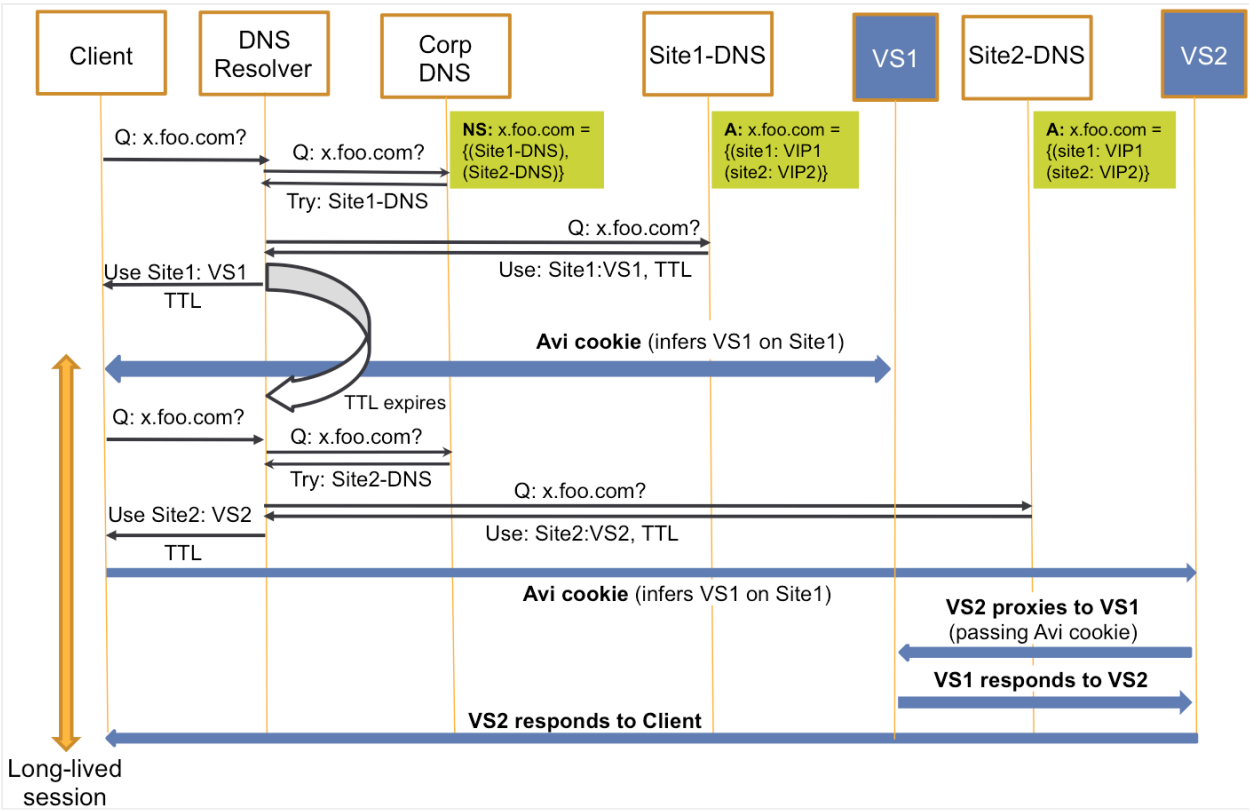

Figure 3. Phase 3 of GSLB cookie-based site persistence

Configuring GSLB Site Cookie Persistence

Outline of Steps to be Taken

The below steps assume a basic GSLB configuration already exists.

1. Configure exactly one federated PKI profile. This one-time operation is most easily done via the Avi UI and applies to all GSLB services. 2. Configure a federated application persistence profile. Multiple such profiles may be defined. 3. Configure a health monitor. Multiple such health monitors may be defined. 4. Configure the GSLB service. Identify it as being site-persistent, and associate it with:
a. One federated application persistence profile
b. One or more health monitors

> *Note: At the time of this writing, all these steps can be accomplished using the Avi UI with the exception of step 4a. Full UI support is planned for a future release.*

Configuration via Avi UI

Step 1. Configure the PKI Profile

- Navigate to Templates > Security > PKI Profile. Click on Create, and be sure to select Is Federated option. This is a one-time operation.

  Note: Where applicable, the `is_federated` option of an Avi object describes its replication scope. If the option is set to false, then the object is visible only within the Controller cluster and its associated Service Engines. If the option is set to true, the object is replicated across the federation.

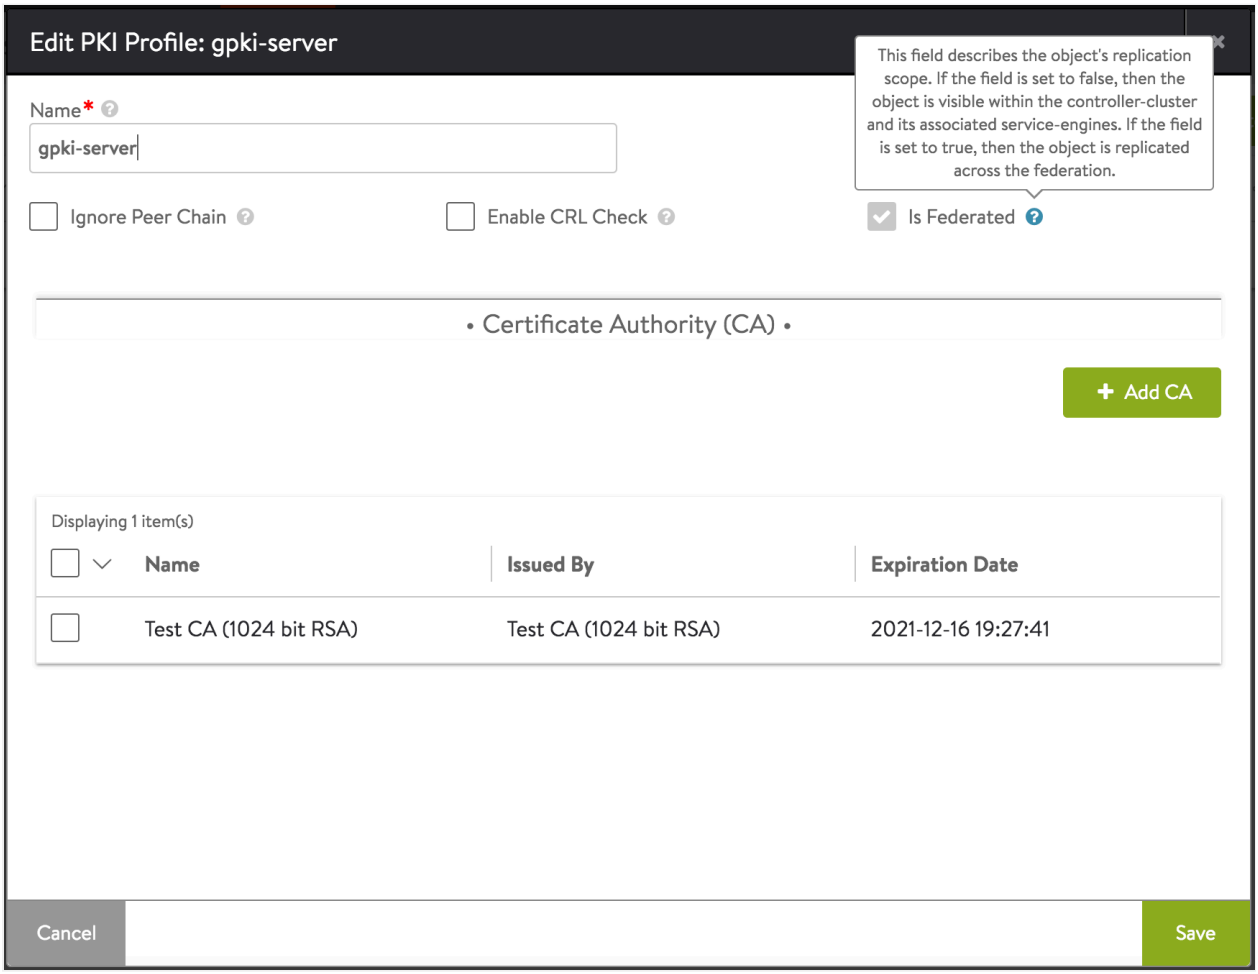

Figure 4. The PKI profile editor Once the federated PKI profile is created and a site-persistent GSLB is enabled, the PKI profile cannot be deleted. Figure 5 shows it listed. If the box at left were checked and the Delete button pressed, the error depicted in figure 6 would appear.

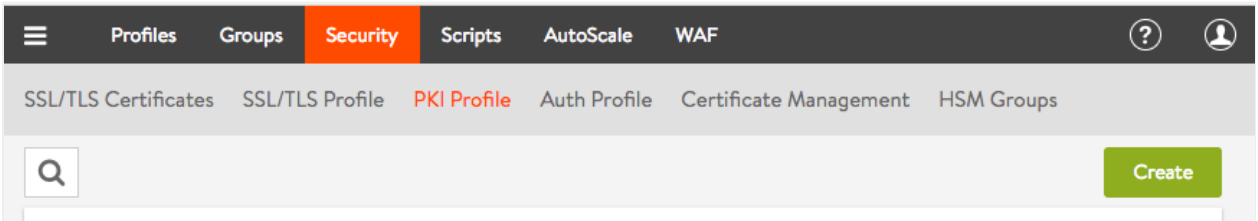

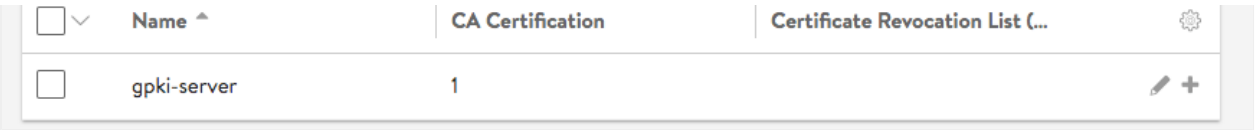

Figure 5. List of PKI profiles

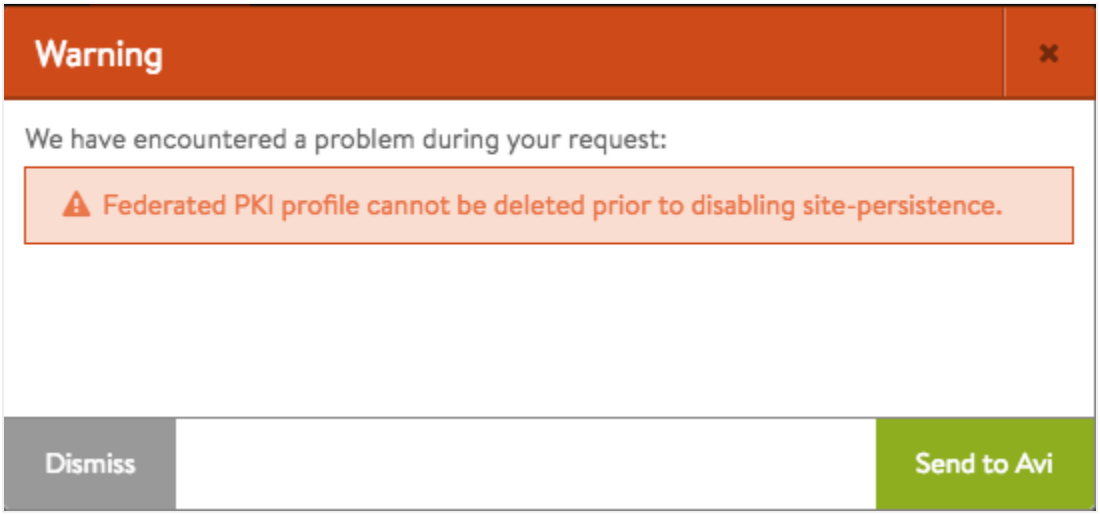

Figure 6. The federated PKI profile can't be deleted if a site-persistent GSLB service is enabled.

Step 2. Configure a Federated Application Persistence Profile

- Navigate to Templates > Profiles > Persistence and click on Create to open the persistence profile editor, as depicted in figure 7. Be sure to set the Type field to GSLB Site and click on the Is Federated option.

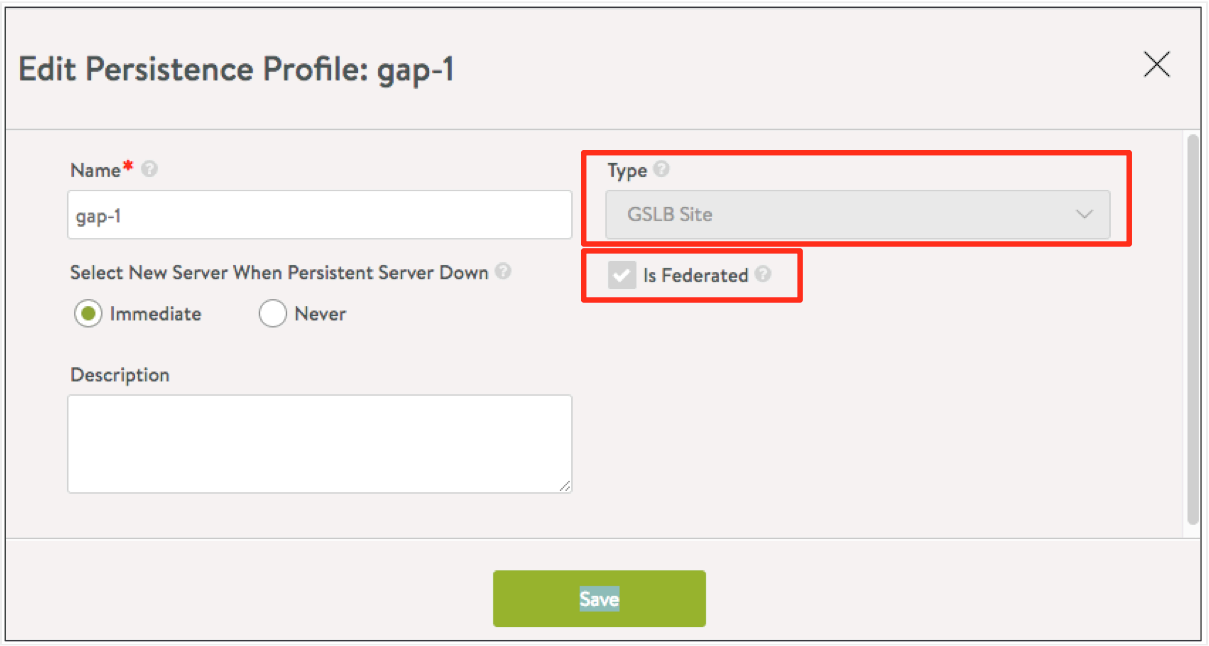

Figure 7. Persistence profile editor

Figure 8 below shows a partial list of configured persistence profiles, the first three of which have their Type field set to GSLB Site.

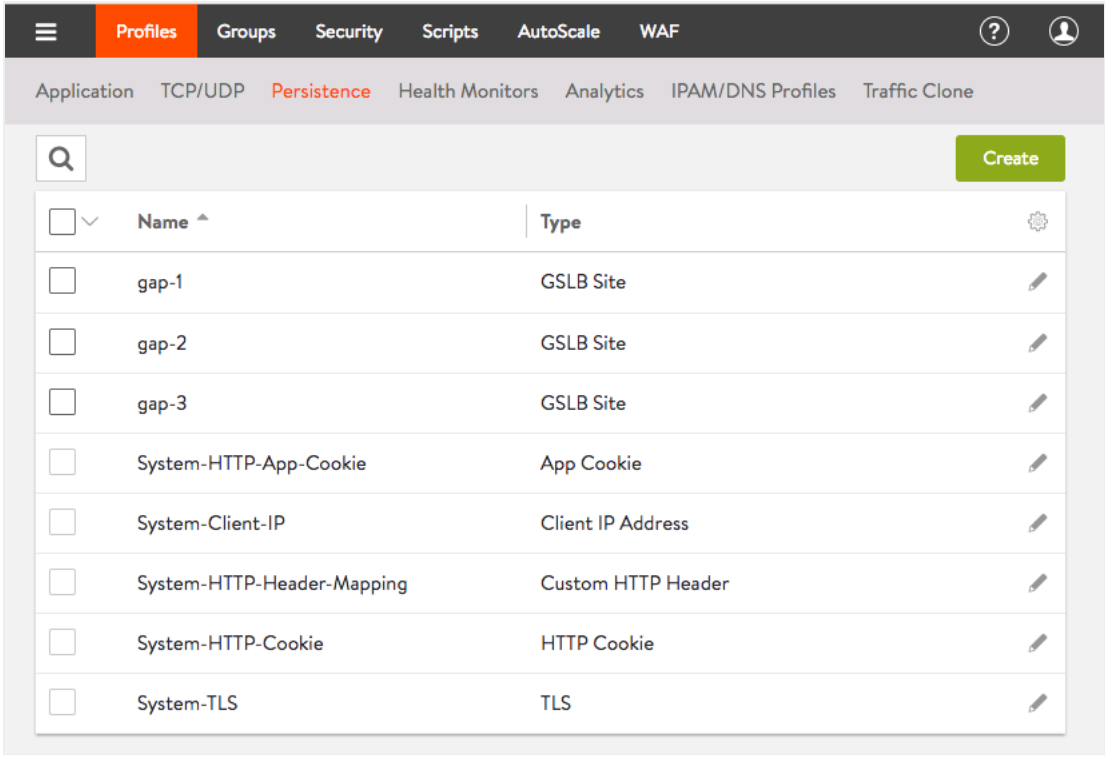

Figure 8. Partial list of configured persistence profiles

Step 3. Configure a Health Monitor

Navigate to Templates > Profiles > Health Monitors and click on Create to open the health monitor editor. Once again, be sure to check the Is Federated option.

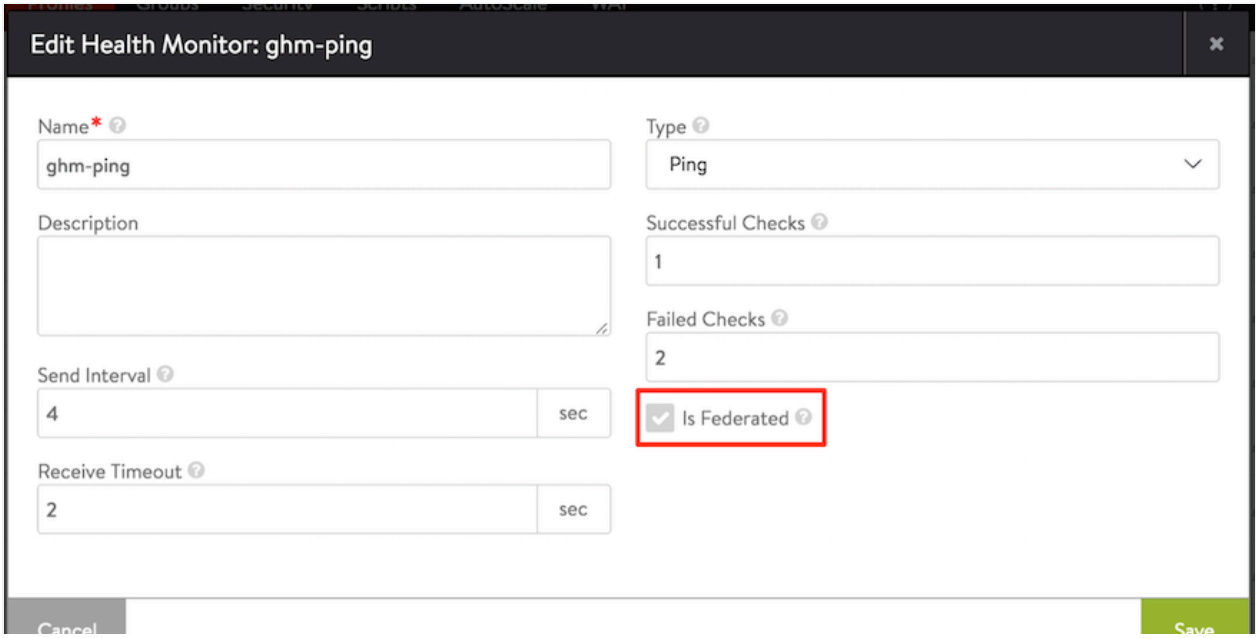

Figure 9. The health monitor editor

Step 4. Configure the GSLB Service

- Navigate to Applications > GSLB Services. Click on Create, and select Advanced Setup. Be sure to specify a health monitor profile and check the Site Persistence option.

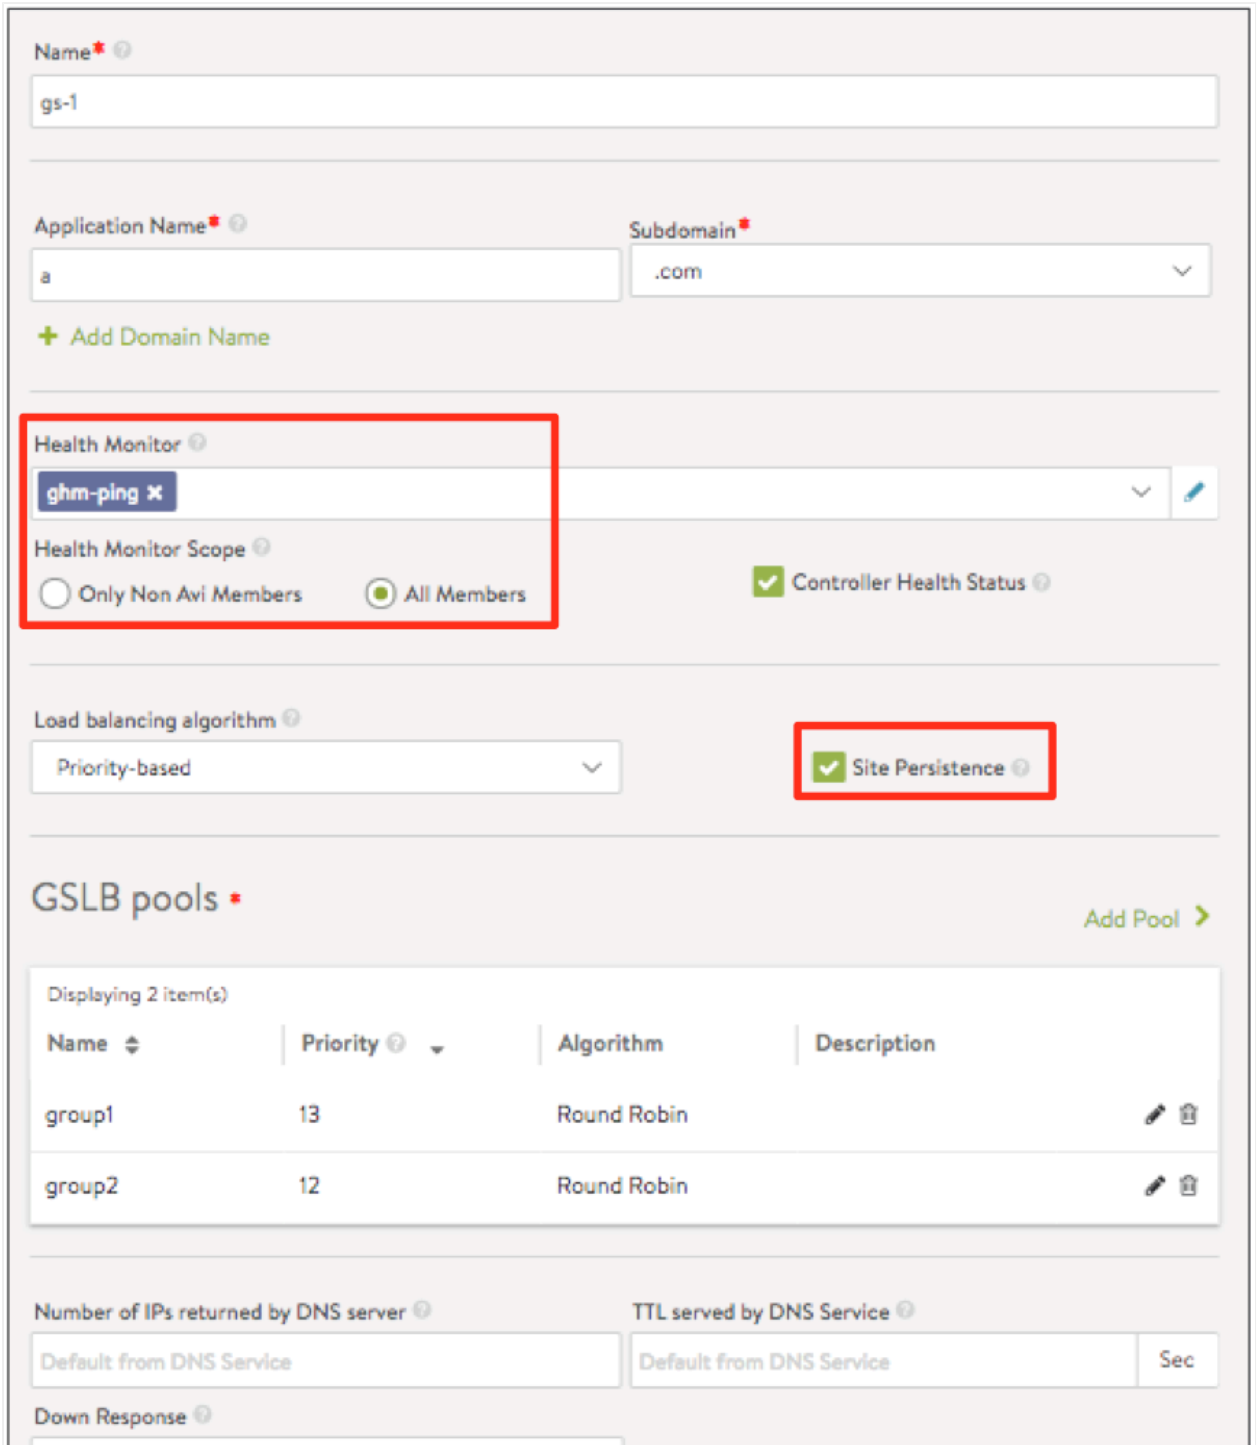

Figure 10. GSLB advanced setup editor

As mentioned, associating the GSLB service with a federated application profile must be performed via the Avi CLI. To sketch the simple steps,

- Log into the Avi `shell` of the appropriate Controller cluster.
- Type
  `configure gslbservice gs-1`
- In response to the `gslbservice` prompt, type
  `application_persistence_profile_ref gap-1`
- To have the association take effect, type
  `save`

For reference, following are all parameters for the enabled GSLB service named `gs-1`:

```
+-----------------------------------+-----------------------------------------------------+
| Field                             | Value                                               |
+-----------------------------------+-----------------------------------------------------+
| uuid                              | gslbservice-2efeea54-12b4-4c1d-9fe0-ffd58e5125c3    |
| name                              | gs-1                                                |
| domain_names[1]                   | a.com                                               |
| groups[1]                         |                                                     |
|   name                            | group1                                              |
|   priority                        | 13                                                  |
|   algorithm                       | GSLB_ALGORITHM_ROUND_ROBIN                          |
|   members[1]                      |                                                     |
|     cluster_uuid                  | cluster-a7ba9c02-adf6-48d7-aa3d-41f664d45f85        |
|     vs_uuid                       | virtualservice-da9efdc9-7204-4b69-afc2-4fccaf961e1d |
|     ip                            | 10.90.173.73                                        |
|     ratio                         | 1                                                   |
|     enabled                       | True                                                |
| groups[2]                         |                                                     |
|   name                            | group2                                              |
|   priority                        | 12                                                  |
|   algorithm                       | GSLB_ALGORITHM_ROUND_ROBIN                          |
|   members[1]                      |                                                     |
|     cluster_uuid                  | cluster-fc6fa719-054d-42d0-a18b-a8a7577a3829        |
|     vs_uuid                       | virtualservice-f4bdb96d-4de3-4b2f-bf51-1bb924783443 |
|     ip                            | 10.90.174.72                                        |
|     ratio                         | 1                                                   |
|     enabled                       | True                                                |
| health_monitor_refs[1]            | ghm-ping                                            |
| controller_health_status_enabled  | True                                                |
| health_monitor_scope              | GSLB_SERVICE_HEALTH_MONITOR_ALL_MEMBERS             |
| enabled                           | True                                                |
| use_edns_client_subnet            | True                                                |
```

```
| wildcard_match                     | False                                              |
| site_persistence_enabled           | True                                               |
| application_persistence_profile_ref | gap-1                                            |
| pool_algorithm                     | GSLB_SERVICE_ALGORITHM_PRIORITY                    |
| min_members                        | 0                                                  |
| is_federated                       | True                                               |
| tenant_ref                         | admin                                              |
+------------------------------------+----------------------------------------------------+
```

Commands for Configuring GSLB Site Cookie Persistence via the Avi CLI

In the below examples, we use the same object names as were used in the above UI configuration, i.e., `gs-1`, `gpki-server`, `gap-1`, and `ghm-ping`. Each shell command has many subcommands; we show only the ones that are especially relevant to GSLB site persistence.

Step 1. Configure the PKI Profile

Shell command: `configure pkiprofile gpki-server`
Subcommands: `is_federated`

Step 2. Configure a Federated Application Persistence Profile

Shell command: `configure applicationpersistenceprofile gap-1`
Subcommands: `is_federated`, `persistence_type`, `server_hm_down_recovery`

Step 3. Configure a Health Monitor

Shell command: `configure healthmonitor ghm-ping`
Subcommands: `is_federated`

Step 4. Configure the GSLB Service

Shell command: `configure gslbservice gs-1`
Subcommands: `application_persistence_profile_ref`, `health_monitor_refs`, `is_federated`, `site_persistence_enabled`

Global Services That Define Both HTTP and HTTPS Ports

Special consideration is required when a global service with site persistence (SP=ON) defines both HTTP and HTTPS ports, be they the default ports (80 and 443) or some other port-pair.

Case 1: Same global application exposes HTTP on port 80 and HTTPS on port 443

You need only set `http_to_https` to True in the application profile associated with the *every* virtual service participating in the global service. In the Avi UI, use the application profile as shown below.

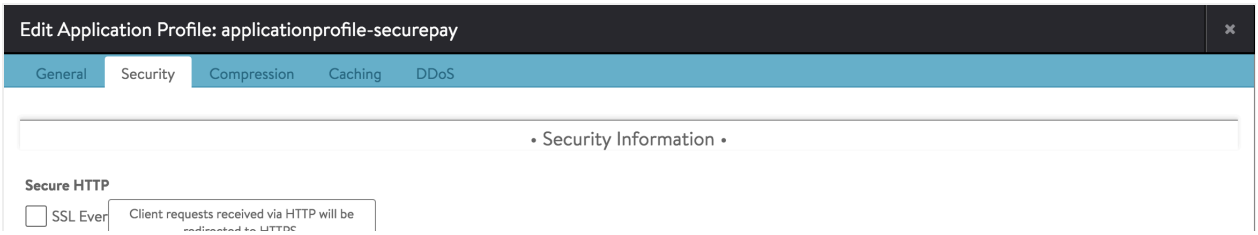

Figure 11. HTTP-to-HTTPS is enabled with a single click.

Case 2: Same global application exposes non-default HTTP and non-default HTTPS ports

To illustrate by example, suppose the virtual services participating in the global service with site persistence (SP=ON) are defined with port 91 for HTTP and port 9443 for HTTPS.

In addition to optioning `http_to_https` ON (via the UI, CLI or REST API), define an HTTP rule for each participating virtual service such that HTTP port 91 is redirected to HTTPS port 9443, as illustrated below.

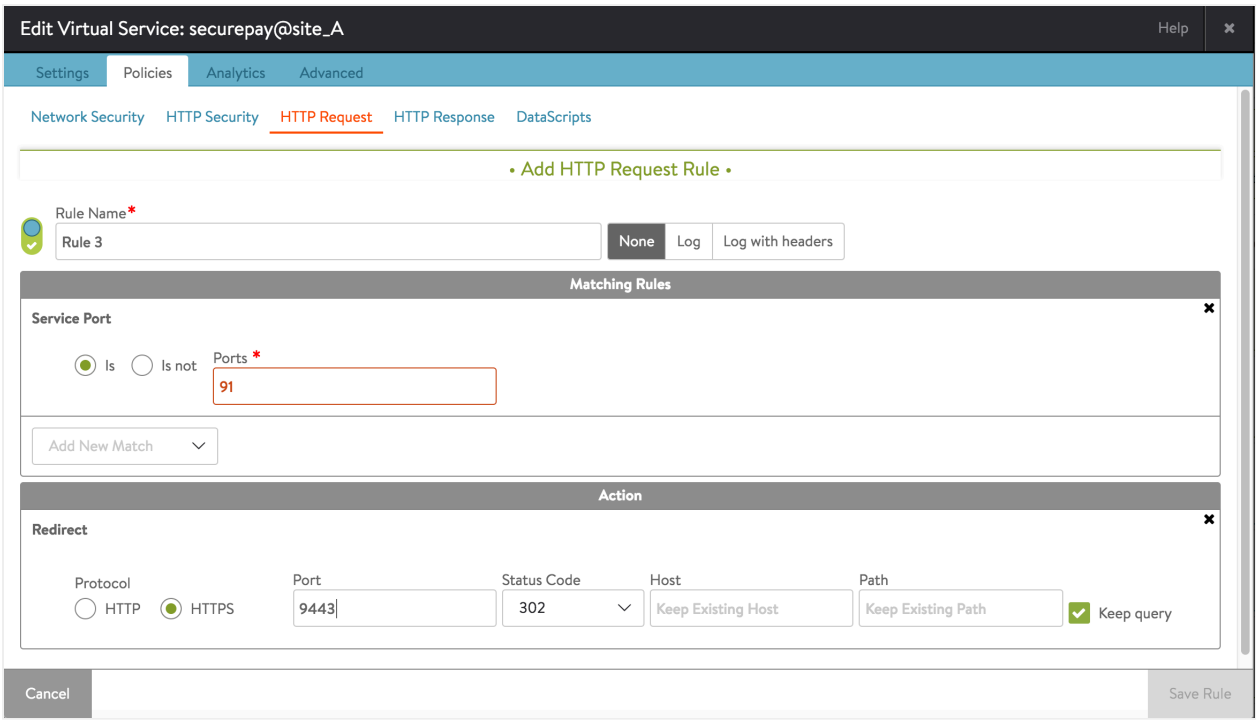

Figure 12. Define an HTTP request rule that defines the ports.

CASE 3: No HTTP-to-HTTPS redirect is in place

Whether the port settings are the default ones (80 and 443) or some other values, without the HTTP-to-HTTPS redirect in place, site-persistence flow will not work.

Operations

This section covers the CLI commands by which to
* Check the operational state of a site-persistent GSLB service. * Determine the percentage of requests that are proxied from other virtual services back to the one to which clients are to be persisted.

In the following CLI sequences we have:

1. A global service named `gs-1`
2. The global service is comprised of two virtual services named `pay@site_A` and `pay@site_B`.
3. These virtual services run on their correspondingly named active sites, `site_A` and `site_B`.
4. A site-persistence proxy pool at each site, correspondingly named `SP-gs-1-pay@site_A` and `SP-gs-1-pay@site_B`. Note that Avi Vantage automatically forms a site's proxy pool name by prepending `SP-` to the hyphenated concatenation of the GSLB service name and the VS name.
5. The operational status with regard to site persistence is up.

GSLB Service Site Persistence Status

The output of the below `show` command reflects points 1 through 5. To the right of the command's output we have inserted annotations to guide you where to look. These data are available from any active site.

> *Note: To view site-persistence-related data you must include the arguments `runtime filter sp_status`.*

```
show gslbservice gs-1 runtime filter sp_status
+------------------------+------------------------------------------------------------------------+
| Field                  | Value                                                                  |
+------------------------+------------------------------------------------------------------------+
| uuid                   | gslbservice-ff1b4e8d-663d-4cb9-932b-d007c81efba6                       |
| name                   | gs-1                                                                   |  POINT 1
| ldr_state              |                                                                        |
|   last_changed_time    | Tue Feb  6 00:11:02 2018 ms(242588) UTC                                |
| flr_state[1]           |                                                                        |
|   status               | SYSERR_SUCCESS                                                         |
|   reason               |                                                                        |
|   site_uuid            | cluster-1e560f44-c898-41c3-818b-3433edbf9391                           |
|   last_changed_time    | Tue Feb  6 00:11:02 2018 ms(904114) UTC                                |
| groups[1]              |                                                                        |
|   name                 | group2                                                                 |
|   members[1]           |                                                                        |
|     cluster_uuid       | cluster-1e560f44-c898-41c3-818b-3433edbf9391                           |
|     site_name          | site_B                                                                 |  POINT 3
|     vs_uuid            | virtualservice-8a68c656-6a89-46d7-b9a5-1b693ae9798a                    |
|     vs_name            | pay@site_B                                                             |  POINT 2
|     ip                 | 10.90.174.72                                                           |
```

```
|     oper_ips[1]         | 10.90.174.72                                                        |
|     vip_type            | AVI_VIP                                                             |
|     services[1]         |                                                                     |
|       port              | 80                                                                  |
|       enable_ssl        | False                                                               |
|       port_range_end    | 80                                                                  |
|     app_type            | APPLICATION_PROFILE_TYPE_HTTP                                       |
|     sp_pools[1]         |                                                                     |
|       uuid              | pool-8a68c656-6a89-46d7-b9a5-1b693ae9798a                           |
|       name              | SP-gs-1-pay@site_B                                                  |  POINT 4
|       num_servers       | 1                                                                   |
|       num_servers_up    | 1                                                                   |
|     controller_status   |                                                                     |
|       state             | OPER_UP                                                             |
|       last_changed_time | Tue Feb  6 00:15:17 2018 ms(352917) UTC                             |
| groups[2]               |                                                                     |
|   name                  | group1                                                              |
|   members[1]            |                                                                     |
|     cluster_uuid        | cluster-3a179b95-dff9-444b-9986-ba89c4e19c44                        |
|     site_name           | site_A                                                              |  POINT 3
|     vs_uuid             | virtualservice-dc871051-35e8-4bec-bd1f-3c63fb6b7087                 |
|     vs_name             | pay@site_A                                                          |  POINT 2
|     ip                  | 10.90.173.73                                                        |
|     oper_ips[1]         | 10.90.173.73                                                        |
|     vip_type            | AVI_VIP                                                             |
|     services[1]         |                                                                     |
|       port              | 80                                                                  |
|       enable_ssl        | False                                                               |
|       port_range_end    | 80                                                                  |
|     app_type            | APPLICATION_PROFILE_TYPE_HTTP                                       |
|     sp_pools[1]         |                                                                     |
|       uuid              | pool-dc871051-35e8-4bec-bd1f-3c63fb6b7087                           |
|       name              | SP-gs-1-pay@site_A                                                  |  POINT 4
|       num_servers       | 1                                                                   |
|       num_servers_up    | 1                                                                   |
|     controller_status   |                                                                     |
|       state             | OPER_UP                                                             |
|       last_changed_time | Tue Feb  6 00:15:17 2018 ms(353741) UTC                             |
| services_state          | Services-In-Sync                                                    |
| tenant_name             | admin                                                               |
| checksum                | e298eb000bb6d5bcaeaaf10d08e609441823c69fc83e7d9a50014769d7ed2b03    |
| sp_oper_status          |                                                                     |
|   state                 | OPER_UP                                                             |  POINT 5
|   last_changed_time     | Tue Feb  6 00:15:17 2018 ms(353976) UTC                             |
+-------------------------+---------------------------------------------------------------------+
```

Status of the GSLB Service's Member Virtual Services

For details about the individual virtual services that comprise a GSLB service, one must log onto the site that pertains. The below `show virtualservice` command was executed on `site_A` to report on a local VS, `pay@site_A`. Note the site-

persistence pool reference toward the very bottom. The SP pool on `site_A` engages the service of some VS on another active site, the site to which the client's request must be persisted. In this example, there's only one other site (`site_B`), but in general there could be many.

```
show virtualservice pay@site_A
+----------------------------------+--------------------------------------------------+
| Field                            | Value                                            |
+----------------------------------+--------------------------------------------------+
| uuid                             | virtualservice-dc871051-35e8-4bec-bd1f-3c63fb6b7087 |
| name                             | pay@site_A                                       |
| enabled                          | True                                             |
| services[1]                      |                                                  |
|   port                           | 80                                               |
|   enable_ssl                     | False                                            |
|   port_range_end                 | 80                                               |
| application_profile_ref          | System-HTTP                                      |
| network_profile_ref              | System-TCP-Proxy                                 |
| pool_ref                         | pay                                              |
| se_group_ref                     | Default-Group                                    |
| analytics_policy                 |                                                  |
|   full_client_logs               |                                                  |
|     enabled                      | True                                             |
|     duration                     | 0 min                                            |
|     all_headers                  | True                                             |
|     throttle                     | 0 per_second                                     |
|   client_insights                | NO_INSIGHTS                                      |
|   udf_log_throttle               | 10 per_second                                    |
|   significant_log_throttle       | 10 per_second                                    |
|   enabled                        | True                                             |
| vrf_context_ref                  | global                                           |
| enable_autogw                    | False                                            |
| analytics_profile_ref            | System-Analytics-Profile                         |
| weight                           | 1                                                |
| delay_fairness                   | False                                            |
| max_cps_per_client               | 0                                                |
| limit_doser                      | False                                            |
| type                             | VS_TYPE_NORMAL                                   |
| cloud_type                       | CLOUD_NONE                                       |
| use_bridge_ip_as_vip             | False                                            |
| flow_dist                        | LOAD_AWARE                                       |
| ign_pool_net_reach               | False                                            |
| ssl_sess_cache_avg_size          | 1024                                             |
| remove_listening_port_on_vs_down | False                                            |
| close_client_conn_on_config_update | False                                          |
| tenant_ref                       | admin                                            |
| cloud_ref                        | Default-Cloud                                    |
| east_west_placement              | False                                            |
| scaleout_ecmp                    | False                                            |
| active_standby_se_tag            | ACTIVE_STANDBY_SE_1                              |
| flow_label_type                  | NO_LABEL                                         |
```

```
| vip[1]                           |                                                    |                      |
|   vip_id                         | 0                                                  |                      |
|   ip_address                     | 10.90.173.73                                       |                      |
|   enabled                        | True                                               |                      |
|   auto_allocate_ip               | False                                              |                      |
|   auto_allocate_floating_ip      | False                                              |                      |
|   avi_allocated_vip              | False                                              |                      |
|   avi_allocated_fip              | False                                              |                      |
| vsvip_ref                        | vsvip-5c8iRv                                       |                      |
| sp_pool_refs[1]                  | SP-gs-1-pay@site_A                                 | SP POOL ON site_A    |
| use_vip_as_snat                  | False                                              |                      |
+----------------------------------+----------------------------------------------------+
```

Proxy Pools Appear Alongside Others

The below `show pool` command, executed on `site_A` illustrates the fact that site-persistence pools appear just as others do. In contrast to the last four listed, the two SP pools have "servers" that are actually virtual services on the one and only other site.

```
show pool
+--------------------------+------+--------------+------------+--------------------+
| Name                     | Port | Cloud        | Oper State | Servers (up/total) |
+--------------------------+------+--------------+------------+--------------------+
| SP-gs-2-securepay@site_A | 80   | Default-Cloud | OPER_UP    | 1/1                |
| SP-gs-1-pay@site_A       | 80   | Default-Cloud | OPER_UP    | 1/1                |
| ship                     | 80   | Default-Cloud | OPER_UP    | 1/1                |
| securepay                | 80   | Default-Cloud | OPER_UP    | 2/2                |
| pay                      | 80   | Default-Cloud | OPER_UP    | 1/1                |
| secureship               | 80   | Default-Cloud | OPER_UP    | 2/2                |
+--------------------------+------+--------------+------------+--------------------+
```

Proxy Pool Status

Details about a proxy pool are not rolled up at the GSLB level. One needs to log onto the site that pertains, and then use the `show pool` command on the proxy pool associated with the particular GSLB service. In the below example, we're logged into `site_A`, looking at the site-persistence pool named `sp-gs-1-pay@site_A`.

Note that the one "server" in the SP pool is identified by the VIP (`10.90.174.72`) of a virtual service on `site_B`.

```
show pool sp-gs-1-pay@site_A
+--------------------------------------+--------------------------------------------------------------------+
| Field                                | Value                                                              |
+--------------------------------------+--------------------------------------------------------------------+
| uuid                                 | pool-dc871051-35e8-4bec-bd1f-3c63fb6b7087                          |
| name                                 | SP-gs-1-pay@site_A                                                 |
| default_server_port                  | 80                                                                 |
| graceful_disable_timeout             | 1 min                                                              |
| connection_ramp_duration             | 10 min                                                             |
| max_concurrent_connections_per_server | 0                                                                  |
```

```
| health_monitor_refs[1]               | ghm-ping                                                                 |
| servers[1]                           |                                                                          |  "SERVER"
|   ip                                 | 10.90.174.72                                                             |  10.90.174
|   hostname                           | 10.90.174.72                                                             |
|   enabled                            | True                                                                     |
|   ratio                              | 1                                                                        |
|   verify_network                     | False                                                                    |
|   resolve_server_by_dns              | False                                                                    |
|   prst_hdr_val                       | 16077db5be5a5402f8185e02769756a3f0deffcdc0ab28fe8a60ac13d0219e32 |
|   static                             | False                                                                    |
|   rewrite_host_header                | False                                                                    |
|   description                        | Gslb site-persistence server                                             |
| server_count                         | 1                                                                        |
| lb_algorithm                         | LB_ALGORITHM_LEAST_CONNECTIONS                                           |
| application_persistence_profile_ref  | gap-1                                                                    |
| inline_health_monitor                | True                                                                     |
| use_service_port                     | True                                                                     |
| capacity_estimation                  | False                                                                    |
| server_auto_scale                    | False                                                                    |
| vrf_ref                              | global                                                                   |
| fewest_tasks_feedback_delay          | 10 sec                                                                   |
| enabled                              | True                                                                     |
| request_queue_enabled                | False                                                                    |
| request_queue_depth                  | 128                                                                      |
| host_check_enabled                   | False                                                                    |
| sni_enabled                          | True                                                                     |
| rewrite_host_header_to_sni           | False                                                                    |
| rewrite_host_header_to_server_name   | False                                                                    |
| lb_algorithm_core_nonaffinity        | 2                                                                        |
| gslb_sp_enabled                      | True                                                                     |
| lookup_server_by_name                | False                                                                    |
| description                          | Gslb site-persistence proxy pool                                         |
| tenant_ref                           | admin                                                                    |
| cloud_ref                            | Default-Cloud                                                            |
+--------------------------------------+--------------------------------------------------------------------------+
```

Determining the Fraction of Client Requests Proxied

On a per-GSLB-service basis, use the Avi UI to monitor the per-pool activity on active sites running the GSLB service's VS members. For each site, collect 1. the in-bound request rate for the GSLB service's local VS, and 2. its SP pool request rate.

Calculate the total for 1 and the total for 2 across all sites. If the overall SP pool rate is large compared to the overall VS request rate, you may wish to increase the value of TTL (see figure 10).