



Application Profile

Avi Technical Reference (v20.1)

Copyright © 2020

Application Profile

[view online](#)

Overview

Application profiles determine the behavior of virtual services, based on application type.

The application profile types and their options are described in the following sections:

- [HTTP Profile](#)
- [DNS Profile](#)
- [Layer 4 Profile](#)
- [SSL Profile](#)
- [Syslog Profile](#)
- [SIP Profile](#)

Dependency on TCP/UDP Profile

The application profile associated with a virtual service may have a dependency on an underlying TCP/UDP profile. For example, an HTTP application profile may be used only if the TCP/UDP profile type used by the virtual service is set to type TCP Proxy. The application profile associated with a virtual service instructs the Service Engine (SE) to proxy the service's application protocol, such as HTTP, and to perform functionality appropriate for that protocol.

Application Profile Tab

Select Templates > Profiles > Applications to open the Application Profiles tab, which includes the following functions:

- Search ? Search against the name of the profile.

- Create ? Opens the Create Application Profile popup.
- Edit ? Opens the Edit Application Profile popup.



- Delete ? Removes an application profile (click its check box) if it is not currently assigned to a virtual service. Note ? If the profile is still associated with any virtual services, the profile cannot be removed. In this case, an error message lists the virtual service that still is referencing the application profile. Neither can any one of the system-standard profiles (as illustrated below) be deleted.

The table on this tab provides the following information for each application profile:

Name ^	Type	
System-DNS	DNS	
System-HTTP	HTTP	
System-L4-Application	L4	
System-SIP	SIP	
System-SSL-Application	L4 SSL/TLS	
System-Secure-HTTP	HTTP	
System-Syslog	SYSLOG	

- **Name ? Name of the Profile.**
- **Type ? Type of application profile, which will be either:**
 - **DNS ? Default for processing DNS traffic.**
 - **HTTP ? Default for processing Layer 7 HTTP traffic.**
 - **L4 ? Catch-all for any virtual service that is not using an application-specific profile.**
 - **L4 SSL/TLS ? Catch-all for any virtual service that is SSL-encrypted and not using an application-specific profile.**
 - **Syslog ? Default for processing Syslog traffic.**
 - **SIP ? Default for processing SIP traffic.**

Note: At the time of this writing, Avi Vantage ships with the templates shown in the above window with the *exception* of a System-SIP template. For it ? and SIP ? to appear in the fourth row above, the template must have been manually created by the user beforehand.

Create/Edit an Application Profile

The Create Application Profile and Edit Application Profile screens share the same interface regardless of the application profile chosen.

The initial settings for a new profile are similar regardless of the type of profile chosen:

- **Name ? Enter a unique name for the profile.**
- **Description ? Enter an optional description for the profile.**

- Type ? Click the appropriate type button to select the application for this profile. Select L4 for none.

HTTP Profile

The HTTP application profile (which is the default) allows Avi Vantage to be a proxy for any HTTP traffic. HTTP-specific functionality such as redirects, content switching, or rewriting server responses to client requests may be applied to a virtual service. The settings apply to all HTTP services that are associated with the HTTP profile. HTTP-specific policies or DataScripts also may be attached directly to a virtual service.

The HTTP profile contains these tabs:

- General
- Security
- Compression
- Caching
- DDoS

HTTP General Tab

The general tab contains HTTP basic settings:

The screenshot shows the 'HTTP Settings' configuration window. It includes the following elements:

- Connection Multiplex ?
- WebSockets Proxy ?
- X-Forwarded-For ?
- Preserve Client IP Address
- XFF Alternate Name ?
- Dropdown menu: X-Forwarded-For

- **Connection Multiplex ?** This option controls the behavior of HTTP 1.0 and 1.1 request switching and server TCP connection reuse. This allows Avi Vantage to reduce the number of open connections maintained by servers and better distribute requests across idle servers, thus reducing server overloading and improving performance for end-users. The exact reduction of connections to servers will depend on how long lived the client connections are, the HTTP version, and how frequently request/responses are utilizing the connection. It is important to understand that "connection" refers to a TCP connection, whereas "request" refers to an HTTP request and subsequent response. HTTP 1.0 and 1.1 allow only a single request/response to go over an open TCP connection at a time. Many browsers attempt to mitigate this bottleneck by opening around six concurrent TCP connections to the destination web site. See Multiplex plus Persistence, below.
- **X-Forwarded-For ?** With this option, Avi Vantage will insert an X-Forwarded-For (XFF) header into the HTTP request headers when the request is passed to the server. The XFF header value contains the original client source IP address. Web servers can use this header for logging client interaction instead of using the layer 3 IP address, which will incorrectly reflect the Service Engine's source NAT address. When enabling this option, the XFF Alternate Name field appears, which allows the XFF header insertion to use a custom HTTP header name. If the XFF header or the custom name supplied already exists in the client request, all instances of that header will first be removed. To add the header without removing pre-existing instances of it, use an HTTP request policy.

- **WebSockets Proxy ?** Enabling WebSockets allows the virtual service to accept a client's Upgrade header request. If the server is listening for WebSockets, the connection between the client and server will be upgraded. WebSocket is a full-duplex TCP protocol. The connection will initially start over HTTP, but once successfully upgraded, all HTTP parsing by Avi Vantage will cease and the connection will be treated as a normal TCP connection.
- **Preserve Client IP Address ?** Clicking this option causes the Avi SE to use the client-IP rather than its own as the source-IP for load-balanced connections from the SE to back-end application servers. Enable IP Routing in the SE group is a prerequisite for enabling this option. Preserve Client IP Address is mutually exclusive with SNAT-ting the virtual services. Connection Multiplexing from HTTP(s) Application Profile cannot be used with Preserve Client IP. Also see
- **Save ?** Select another tab from the top menu to continue editing or Save to return to the Application Profiles tab. See also the [Preserve Client IP](#) article.

Multiplex plus Persistence

Multiplexing behavior changes with server persistence enabled:

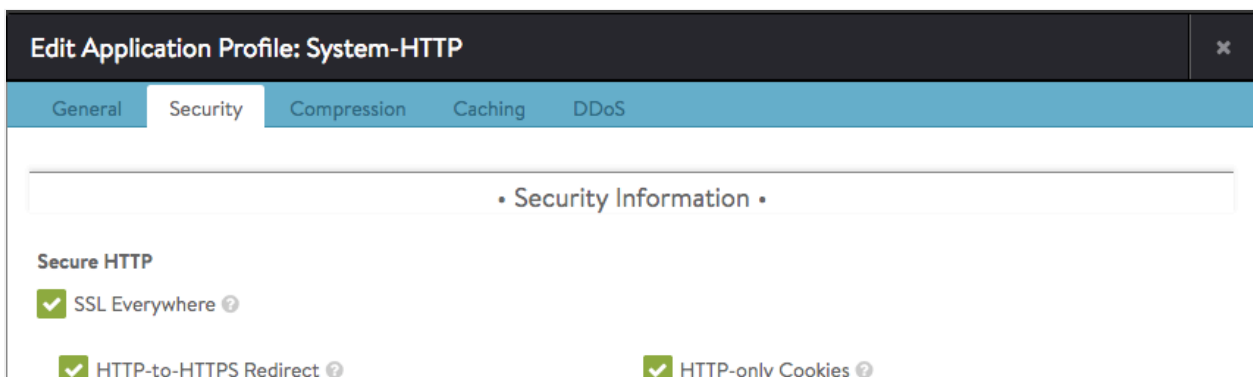
- **Multiplex enabled, Persistence disabled ?** Client connections and their requests are decoupled from the server side of the Service Engine. Requests are load-balanced across the servers in the pool using either new or pre-existing connections to those servers. The connections to the servers may be shared by requests from any clients.
- **Multiplex enabled, Persistence enabled ?** Client connections and their requests are sent to a single server. These requests may share connections with other clients who are persisted to the same server. Load balancing of HTTP requests is not performed.
- **Multiplex disabled, Persistence enabled ?** Avi Vantage opens a new TCP connection to the server for each connection received from the client. Connections are not shared with other clients. All requests received through all connections from the same client are sent to one server. HTTP client browsers may open many concurrent connections, and the number of client connections will be the same as the number of server connections.
- **Multiplex disabled, Persistence disabled ?** Connections between the client and server are one-to-one. Requests remain on the same connection they began on. Multiple connections from the same client may be distributed among the available servers.

HTTP Security

The Security tab of the HTTP application profile controls the security settings for HTTP applications that are associated with the profile:

Security Information

The HTTP security settings affect how a virtual service should handle HTTPS. If a virtual service is configured only for HTTP, any HTTPS settings in this section will not apply. Only if the virtual service is configured for HTTPS, or HTTP and HTTPS, will the settings take effect.



Secure Cookies ?
 Rewrite Server Redirects to HTTPS ?

HTTP Strict Transport Security (HSTS) ?
 X-Forwarded-Proto ?

Duration ?

↳ 365 days

↳ includeSubdomains ?

• Client SSL Certificate Validation •

Validation Type ? None **Request** Required

PKI Profile ?

Select PKI Profile ▼

Add HTTP Request Headers ?

HTTP Header Name	HTTP Header Value
Header Name	Header Value ▼ +

Cancel Save

More granular settings also may be configured using [policies](#) or [DataScripts](#).

- **SSL Everywhere ?** This option enables all of the following options, which together provide the recommended security for HTTPS traffic.
- **HTTP to HTTPS Redirect ?** For a single virtual service configured with both an HTTP service port (SSL disabled) and an HTTPS service port (SSL enabled), this feature will automatically redirect clients from the insecure to the secure port. For instance, clients who type `www.avinetworks.com` into their browser will automatically be redirected to `https://www.avinetworks.com`. If the virtual service does not have both an HTTP and HTTPS service port configured, this feature will not activate. For two virtual services (one with HTTP and another on the same IP address listening to HTTPS), an HTTP request policy must be created to manually redirect the protocol and port.
- **Secure Cookies ?** When Avi Vantage is serving as an SSL proxy for the back-end servers in the virtual service's pool, Avi Vantage communicates with the client over SSL. However, if Avi Vantage communicates with the backend servers over HTTP (not over SSL), the servers will incorrectly return responses as HTTP. As a result, cookies that should be marked as secure will not be so marked. Enabling secure cookies will mark any server cookies with the Secure flag, which tells clients to send only this cookie to the virtual service over HTTPS. This feature will only activate when applied to a virtual service with SSL/TLS termination enabled.
- **HTTP Strict Transport Security (HSTS) ?** Strict Transport Security uses a header to inform client browsers that this site should be accessed only over SSL/TLS. The HSTS header is sent in all HTTP responses, including error responses. This feature mitigates man-in-the-middle attacks that can force a client's secure SSL/TLS session to connect through insecure HTTP. HSTS has a Duration setting that tells clients the SSL/TLS preference should remain in effect for the specified number of days. Starting with release 18.2.2, users can also use the Avi UI to enable the insertion of the `includeSubdomains` directive in the HSTS header. Doing so signals the user agent that the HSTS policy applies to this HSTS host as well as any subdomains of the host's domain name. This setting will activate only on a virtual service

that is configured to terminate SSL/TLS. Note: If a virtual service is set temporarily to support SSL/TLS and HSTS has been set, it cannot gracefully be downgraded back to HTTP. Client browsers will refuse to accept the site over HTTP. When HSTS is in effect, clients will not accept a self-signed certificate.

- **HTTP-only Cookies ?** This marks server cookies as HTTPOnly, which means the cookies cannot be viewed or used by third parties, including Javascript or other web sites. This feature will activate for any HTTP or terminated HTTPS virtual service.
- **Rewrite Server Redirects to HTTPS ?** When a virtual service terminates client SSL/TLS and then passes requests to the server as HTTP, many servers assume that the connection to the client is HTTP. Absolute redirects generated by the server may therefore include the protocol, such as `http://www.avinetworks.com`. If the server returns a redirect with HTTP in the location header, this feature will rewrite it to HTTPS. Also, if the server returns a redirect for its own IP address, this will be rewritten to the hostname requested by the client. If the server returns redirects for host names other than what the client requested, they will not be altered. Note: Consider creating an HTTP response policy if greater granularity is required when rewriting redirects. This feature will activate only if the virtual service has both HTTP and HTTPS service ports configured.
- **X-Forwarded-Proto ?** Enabling this option makes Avi Vantage insert the X-Forwarded-Proto header into HTTP requests sent to the server, which informs that server whether the client connected to Avi Vantage over HTTP or HTTPS. This feature activates for any HTTP or HTTPS virtual service.


Client SSL Certificate Validation

Avi Vantage can validate the certificates presented by clients, by checking them against a Client Revocation List (CRL). Further options allow passing certificate information to the server through HTTP headers.

- **Validation Type ?** Enables client validation based on their SSL certificates.
 - **None ?** Disables validation of client certificates.
 - **Request ?** This setting expects clients to present a client certificate. If a client does not present a certificate, or if the certificate fails the CRL check, the client connection and requests are still forwarded to the destination server. This allows Avi Vantage to forward the client's certificate to the server in an HTTP header, so that the server may make the final determination to allow or deny the client.
 - **Require ?** Avi Vantage requires a certificate to be presented by the client, and the certificate must pass the CRL check. The client certificate, or relevant fields, may still be passed to the server through an HTTP header.
- **PKI Profile ?** The Public Key Infrastructure (PKI) profile contains configured certificate authority (CA) and the CRL. A PKI profile is not necessary if validation is set to Request, but is required if validation is set to Require.
- **HTTP Header Name ?** Optionally, Avi Vantage may insert the client's certificate, or parts of it, into a new HTTP header to be sent to the server. To insert a header, this field is used to determine the name of the header.
- **HTTP Header Value ?** Used with the HTTP Header Name field, the Value field is used to determine the portion of the client certificate to insert into the HTTP header sent to the server. Using the plus icon, additional headers may be inserted. This action may be in addition to any performed by HTTP policies or DataScripts, which could also be used to insert headers in requests sent to the destination servers.

Compression

The Compression tab permits one to view or edit the application profile's compression settings:



Enable Compression
 Compression Mode [?]
 Auto Custom
 Compressible Content Types [?]
 System-Compressible-Content-Types
 Remove Accept Encoding Header [?]
 Cancel Save

The compression option enables HTTP Gzip compression for responses from Avi Vantage to the client. Compression is an HTTP 1.1 standard for reducing the size of text-based data using the Gzip algorithm. The typical compression ratio for HTML, Javascript, CSS, and similar text content types is about 75%, meaning that a 20-KB file may be compressed to 5 KB before being sent across the Internet, thus reducing the transmission time by a similar percentage.

The compression percentage achieved can be viewed using the Client Logs tab of the virtual service. This may require enabling full client logs on the virtual service's Analytics tab to log some or all client requests. The logs will include a field showing the compression percentage with each HTTP response.

Note: It is highly recommended to enable compression in conjunction with caching, which together can dramatically reduce the CPU costs of compressing content. When both compression and caching are enabled, an object such as the index.html file will need to be compressed only one time. After an object is compressed, the compressed object is served out of the cache for subsequent requests. Avi Vantage does not needlessly re-compress the object for every client request. For clients that do not support compression, Avi Vantage also will cache an uncompressed version of the object.

To specify compression settings:

- Check the Compression checkbox to enable compression. You may only change compression settings after enabling this feature.
- Select either Auto or Custom, which enables different levels of compression for different clients. For instance, filters can be created to provide aggressive compression levels for slow mobile clients while disabling compression for fast clients from the local intranet. Auto is recommended, to dynamically tune the settings based on clients and available Service Engine CPU resources.
 - Auto mode enables Avi Vantage to determine the optimal settings.

Note: By default, the Compression Mode is Auto. The content compression depends on the client's RTT, as mentioned below:

```

* RTT less than 10ms, no compression
* RTT 10 to 200ms, normal compression
* RTT above 200ms, aggressive compression
  
```

- Custom mode allows creation of custom filters that provide more granular control over who should receive what level of compression.
- Compressible Content Types determine which HTTP Content-Types are eligible to be compressed. This field points to a String Group which contains the compressible type list.
- Remove Accept Encoding Header removes the Accept Encoding header, which is sent by HTTP 1.1 clients to indicate they are able to accept compressed content. Removing the header from the request prior to sending the request to

the server allows Avi Vantage to ensure the server will not compress the responses. Only Avi Vantage will perform compression.

Custom Compression

To create a custom compression filter:

The screenshot shows the 'New Application Profile' configuration window with the 'Compression' tab selected. The 'Enable Compression' checkbox is checked. The 'Compression Mode' is set to 'Custom'. The 'Compressible Content Types' dropdown is set to 'System-Compressible-Content-Types'. The 'Remove Accept Encoding Header' checkbox is checked. Below this, there is a section for 'Add Compression Filter' with a filter named 'Filter 1'. The 'Matching Rules' section includes 'Client IP Address' with 'Is in' selected and a 'Select IP Group' dropdown, and 'User Agent contains' with a 'Select String Group' dropdown. The 'Action' section has 'Compression' set to 'Normal'. At the bottom, there are 'Cancel' and 'Save Filter' buttons.

1. Click Add New Filter to create a custom filter.

2. Enter the following:

- Filter Name ? Provide a unique name for the filter (optional).
- Matching Rules ? determine if the client (via Client IP or User Agent string) is eligible to be compressed via the associated Action. If both Client IP and User Agent rules are populated, then both must be true for the compression action to fire.
 - Client IP Address allows you to use an IP Group to specify eligible client IP addresses. For example, an IP Group called Intranet that contains a list of all internal IP address ranges. Clearing the Is In button reverses this logic, meaning that any client that is not coming from an internal IP network will match the filter.
 - User Agent matches the client's User Agent string against an eligible list contained within a String Group. The User Agent is a header presented by clients indicating the type of browser or device they may be using. The System-Devices-Mobile Group contains a list of HTTP User Agent strings for common mobile browsers.

3. The Action section determines what will happen to clients or requests that meet the Match criteria, specifically the level of HTTP compression that will be used.

- Aggressive compression uses Gzip level 6, which will compress text content by about 80% while requiring more CPU resources from both Avi Vantage and the client.
- Normal compression uses Gzip level 1, which will compress text content by about 75%, which provides a good mix between compression ratio and the CPU resources consumed by both Avi Vantage and the client.
- No Compression disables compression. For clients coming from very fast, high bandwidth and low latency connections, such as within the same data center, compression may actually slow down the transmission time and consume unnecessary CPU resources.

HTTP Caching

Avi Vantage can cache HTTP content, thereby enabling faster page load times for clients and reduced workloads for both servers and Avi Vantage. When a server sends a response, such as logo.jpg, Avi Vantage can add the object to its cache and serve it to subsequent clients that request the same object. This can reduce the number of connections and requests sent to the server.

Enabling caching and compression allows Avi Vantage to compress text-based objects and store both the compressed and original uncompressed versions in the cache. Subsequent requests from clients that support compression will be served from the cache, meaning that Avi Vantage will need not compress every object every time, which greatly reduces the compression workload.

Note: Regardless of the configured caching policy, an object can be cached only if it is [eligible for caching](#). Some objects may not be eligible for caching.

By default, caching is and appears off, as shown at right. Click the box to enable caching.

The screenshot shows the 'New Application Profile' configuration window with the 'Caching' tab selected. The 'Caching' section is expanded, showing the following settings:

- Enable Caching
- X-Cache ?
- Age Header ?
- Date Header ?
- Cacheable Object Size* ?
 - Min. bytes: 100
 - Max. bytes: 4194304
- Cache Expire Time* ?
 - 600 min
- Heuristic Expire ?
- Cache URI with Query Arguments ?
- Cacheable MIME Types ?
 - Enter String or Select String Group
- Non-Cacheable MIME Types ?
 - Enter String or Select String Group

Buttons for 'Add Group' are located below the MIME Type dropdowns. At the bottom of the window are 'Cancel' and 'Save' buttons.

The following parameters all are optional:

- X-Cache ? Avi Vantage will add an HTTP header labeled X-Cache for any response sent to the client that was served from the cache. This header is informational only, and will indicate the object was served from an intermediary cache.

- **Age Header ?** Avi Vantage will add a header to the content served from cache that indicates to the client the number of seconds that the object has been in an intermediate cache. For example, if the originating server declared that the object should expire after 10 minutes and it has been in the Avi Vantage cache for 5 minutes, then the client will know that it should only cache the object locally for 5 more minutes.
- **Date Header ?** If a date header was not added by the server, then Avi Vantage will add a date header to the object served from its HTTP cache. This header indicates to the client when the object was originally sent by the server to the HTTP cache in Avi Vantage.
- **Cacheable Object Size ?** The minimum and maximum size of an object (image, script, and so on) that can be stored in the Avi Vantage HTTP cache, in bytes. Most objects smaller than 100 bytes are web beacons and should not be cached despite being image objects.
- **Cache Expire Time ?** An intermediate cache must be able to guarantee that it is not serving stale content. If the server sends headers indicating how long the content can be cached (such as cache control), then Avi Vantage will use those values. If the server does not send expiration timeouts and Avi Vantage is unable to make a strong determination of freshness, then Avi Vantage will store the object for no longer than the duration of time specified by the Cache Expire Time.
- **Heuristic Expire ?** If a response object from the server does not include the Cache-Control header but does include an If-Modified-Since header, then Avi Vantage will use this time to calculate the cache-control expiration, which will supersede the Cache Expire Time setting for this object.
- **Cache URL with Query Arguments ?** This option allows caching of objects whose URI includes a query argument. Disabling this option prevents caching these objects. When enabled, the request must match the URI query to be considered a hit. Below are two examples of URIs that include queries. The first example may be a legitimate use case for caching a generic search, while the second may be a unique request posing a security liability to the cache.
 - www.search.com/search.asp?search=caching
 - www.foo.com/index.html?loginID=User
- **Cacheable MIME Types ?** Statically defines a list of cacheable objects. This may be a string group, such as System-Cacheable-Resource-Types, or a custom comma-separated list of MIME types that Avi Vantage should cache. If no MIME types are listed in this field, then Avi Vantage will by default assume that any object is eligible for caching.
- **Non-Cacheable MIME Types ?** Statically define a list of objects that are not cacheable. This creates a blacklist that is the opposite of the cacheable list.

HTTP DDoS

The Distributed Denial of Service (DDoS) section allows configuration of mitigation controls for HTTP and the underlying TCP protocols. By default, Avi Vantage is configured to protect itself from a number of types of attacks. For instance, if a virtual service is targeted by a SYN flood attack, Avi Vantage will activate SYN cookies to validate clients before opening connections. Many of the options listed below are not quite as straightforward, as bursts of data may be normal for the application. Avi Vantage provides a number of knobs to modify the default behavior to ensure optimal protection.

In addition to the DDoS settings described below, Avi Vantage also can implement connection limits to a virtual service and a pool, configured through the Advanced properties page. Virtual services also may be configured with connection rate limits and burst limits in the Network Security Policies section. Because these settings apply on to an individual virtual service and pool, they are not configured within the profile.

The screenshot shows the configuration interface for the DDoS section. The tabs at the top are General, Security, Compression, Caching, and DDoS. The main content area is titled "• HTTP Limit Settings •" and contains two sections: "HTTP Timeout Settings" and "HTTP Size Settings".

HTTP Timeout Settings		HTTP Size Settings	
Client Header Timeout ?	Client Body Timeout ?	Client Post Body Size ?	Client Header Size ?
10000 ms	30000 ms	0 KB	12 KB
HTTP Keep-Alive Timeout ?	Post Accept Timeout ?	Client Request Size ?	
00000	00000		

30000	ms	30000	ms	48	KB
<input type="checkbox"/> Send Keep-Alive header ?					
<input type="checkbox"/> Use App Keep-Alive Timeout ?					
• Rate Limit HTTP and TCP Settings •					
Rate Limit Connections from a Client ?					
Threshold ?	Time Period *	Action *			
0	1 sec	Report Only			
Add a Rate Limit					
<input type="text"/>					
Cancel				Save	

HTTP Limits

The first step in mitigating HTTP-based denial of service attacks is to set parameters for the transfer of headers and requests from clients. Many of these settings protect against variations of HTTP SlowLoris and SlowPOST attacks, in which a client opens a valid connection then very slowly streams the request headers or POSTs a file. This type of attack is intended to overwhelm the server (in this case the Service Engine) by tying up buffers and connections. Clients that exceed the limits defined below will have that TCP connection reset and a log generated. This does not prevent the client from initiating a new connection and does not interrupt other connections the same client may have open.

- **Client Header Timeout ?** Set the maximum length of time the client is allowed for successfully transmitting the complete headers of a request. The default is 10 seconds.
- **HTTP Keep-alive Timeout ?** Set the maximum length of time an HTTP 1.0 or 1.1 connection may be idle. This affects only client-to-Vantage interaction. The Avi Vantage-to-server keep-alive is governed through the Connection Multiplex feature.
- **Client Body Timeout ?** Set the maximum length of time for the client to send a message body. This usually affects only clients that are POSTing (uploading) objects. The default value of 0 disables this timeout.
- **Post Accept Timeout ?** Once a TCP three-way handshake has successfully completed, the client has this much time to send the first byte of the request header. Once the first byte has been received, this timer is satisfied and the client header timeout (described above) kicks in.
- **Send Keep-Alive header ?** Check this to send the HTTP keep-alive header to the client.
- **Use App Keep-Alive Timeout ?** When the above parameter is checked such that keep-alive headers are sent to the client, a timeout value needs to be specified therein. If this box is unchecked, Avi Vantage will use the value specified in the HTTP Keep-Alive Timeout field. If it is checked, the timeout sent by the application will be honored.
- **Client Post Body Size ?** Set the maximum size of the body of a client request. This generally limits the size of a client POST. Setting this value to 0 disables this size limit.
- **Client Request Size ?** Set the maximum combined size of all the headers in a client request.
- **Client Header Size ?** Set the maximum size of a single header in a client request.

Rate Limits

This section controls the rate at which clients may interact with the site. Each enabled rate limit has three settings:

- **Threshold ?** The client has violated the rate limit when the defined threshold of connections, packets, or HTTP requests have occurred within the specified time Period.
- **Time Period ?** The client has violated the rate limit when the defined threshold of connections, packets, or HTTP requests have occurred within the specified time Period.

- Action ? Select the action to perform when a client has exceeded the rate limit. The options will depend on whether the limit is a TCP limit or an HTTP limit.
 - Report Only ? A log is generated on the virtual server log page. By default, no action is taken. However, this option may be used with an alert to generate an alert action to send a notice to a remote destination or to take action through a ControlScript.
 - Drop SYN Packets ? For TCP-based limits, silently discard TCP SYNs from the client. Avi Vantage also will generate a log. However, during high volumes of DoS traffic, repetitive logs may be skipped.
 - Send TCP RST ? Reset client TCP connection attempts. While more graceful than the Drop SYN Packet option, sending a TCP reset does generate extra packets for the reset, versus the Drop SYN Packet option which does not send a client response. Avi Vantage also will generate a log. However, during high volumes of DoS traffic, repetitive logs may be skipped.
 - Close TCP Connection ? Resets a client TCP connection for an HTTP rate limit violation.
 - Send HTTP Local Response ? The Service Engine will send an HTTP response directly to the client without forwarding the request on to the server. Select the HTTP status code of the response, and optionally a response page.
 - Send HTTP Redirect ? Redirect the client to another location.

The following rate limits may be configured.

- Rate Limit Connections from a Client ? Rate limit all connections made from any single client IP address to the virtual service.
- Rate Limit Requests from a Client to all URLs ? Rate limit all HTTP requests from any single client IP address to all URLs of the virtual service.
- Rate Limit Requests from all Clients to a URL ? Rate limit all HTTP requests from all client IP addresses to any single URL.
- Rate Limit Requests from a Client to a URL ? Rate limit all HTTP requests from any single client IP address to any single URL.
- Rate Limit Failed Requests from a Client to all URLs ? Rate limit all requests from a client for a specified period of time once the count of failed requests from that client crosses a threshold for that period. Clients are tracked based on their IP address. Requests are deemed failed based on client or server side error status codes, consistent with how Avi Vantage logs and how metrics subsystems mark failed requests.
- Rate Limit Failed Requests from all Clients to a URL ? Rate limit all requests to a URI for a specified period of time once the count of failed requests to that URI crosses a threshold for that period. Requests are deemed failed based on client- or server-side error status codes, consistent with how Avi Vantage logs and metrics subsystems mark failed requests.
- Rate Limit Failed Requests from a Client to a URL ? Rate limit all requests from a client to a URI for a specified period of time once the count of failed requests from that client to the URI crosses a threshold for that period. Requests are deemed failed based on client- or server-side error status codes, consistent with how Avi Vantage logs and metrics subsystems mark failed requests.
- Rate Limit Scans from a Client to all URLs ? Automatically track clients and classify them into three groups: Good, Bad, and Unknown. Clients are tracked based on their IP address. Clients are added to the Good group when the Avi Vantage scan detection system builds history of requests from the clients that complete successfully. Clients are added to the Unknown group when there is insufficient history about them. Clients with a history of failed requests are added to the Bad group and their requests are rate limited with stricter thresholds than the Unknown clients group. The Avi Vantage scan detection system automatically tunes itself so that the Good, Bad, and Unknown client-IP group members change dynamically with changes in traffic patterns through Avi Vantage. In other words, if a change to the website causes mass failures (such as 404 errors) for most customers, Avi Vantage adapts and does not mark all clients as attempting to scan the site.
- Rate Limit Scans from all Clients to all URLs ? Similar to the previous limit, but restricts the scanning from all clients as a single entity rather than individually. Once a limit is collectively reached by all clients, any client that sends the next failed request will be reset.

DNS Profile

A DNS application profile specifies settings dictating Avi Vantage's request-response handling. By default, this profile will set the virtual service's port number to 53, and the network protocol to UDP with per-packet parsing.

New Application Profile:

General

Name* Type L4 SSL/TLS L4 DNS SYSLOG HTTP SIP

Description

• DNS Settings •

Number of IPs returned by DNS server Negative TTL Sec

TTL Sec (Options for) Invalid DNS Query processing DNS_ERROR_RESPONSE_NONE ▼

Subnet prefix length Respond to AAAA queries with empty response

Process EDNS Extensions

• DNS Request Rate Limiter Settings •

Rate Limit Connections from a Client

Threshold Time Period Seconds Action Action ▼

• Advanced Settings •

Preserve Client IP Address

Valid subdomains

Authoritative Domain Names

Cancel Save

- **Number of IPs returned by DNS server ?** Specifies the number of IP addresses returned by the DNS service. Default is 1. Enter 0 to return all IP addresses. Otherwise, the valid range is 1 to 20.
- **TTL ?** The time in seconds (default = 30) a served DNS response is to be considered valid by requestors of the DNS service. Valid range is 1 to 86400 seconds.
- **Subnet prefix length ?** This length is used in concert with the DNS client subnet (ECS) option. When the incoming request does not have any ECS and the prefix length is specified, Avi Vantage inserts an ECS option in the request to upstream servers. Valid lengths range from 1 to 32.
- **Process EDNS Extensions ?** This option makes the DNS service aware of the [Extension mechanism for DNS \(EDNS\)](#). EDNS extensions are parsed and shown in logs. For GSLB services, the EDNS subnet option can be used to influence load balancing. EDNS support was added in Avi Vantage 17.1.3.

- **Negative TTL ?** Specifies the TTL value (in seconds) for SOA (Start of Authority) (corresponding to a authoritative domain owned by this DNS Virtual Service) record's minimum TTL served by the DNS Virtual Service. Negative TTL is a value in the range 0-86400.
- **(Options for) Invalid DNS Query processing ?** Specifies whether the DNS service should drop or respond to a client when processing its request results in an error. By default, such a request is dropped without any response, or passed through to a passthrough pool, if configured. When set to respond, an appropriate response is sent to the client, e.g., NXDOMAIN response for non-existent records, empty NOERROR response for unsupported queries, et cetera.
- **Respond to AAAA queries with empty response ?** Enable this option to have the DNS service respond to AAAA queries with an empty response when there are only IPv4 records.
- **Rate Limit Connections from a Client ?** Limits connections made from any single client IP address to the DNS virtual service for which this profile applies. The default (=0) is interpreted as no rate limiting.
- **Threshold ?** Specifies the maximum number of connections or requests or packets that will be processed in the time value specified in the Time Period field (legitimate values range from 10 to 2500). A higher number will result in rate limiting. Specifying a number higher than 0 makes the Time Period field mandatory.
- **Time Period ?** The span of time, in seconds, during which Avi Vantage monitors for exceeded threshold. The allowed range is from 1 to 300. Avi Vantage calculates and takes specified action, if the inbound request rate is exceeded. This rate is the ratio of maximum number to the time span.
- **Action ?** Choose one of three actions from the pulldown to be performed when rate limiting is required: Report Only, Drop SYN Packets, or Send TCP RST.
- **Preserve Client IP Address ?** Click this option ON to have the client IP address pass through to the back end. Be sure you understand what the back-end DNS servers expect and what they will do when offered the client IP address. This option is not compatible with connection multiplexing.
- **Valid subdomains ?** A comma-delimited whitelist of subdomain names. Identifies the subdomains serviced by the DNS virtual service with which this profile is associated; all others will not be processed. This option's best use is in the context of GSLB, in which the GSLB DNS' sole purpose is to return IP addresses corresponding to the global applications being served. Valid subdomains are configured with ends-with semantics.
- **Authoritative Domain Names ?** A comma-delimited set of domain names for which the GSLB DNS' SEs can provide authoritative translation of FQDNs to IP addresses. Queries for FQDNs that are subdomains of these domains and do not have any DNS record in Avi are either dropped or an NXDOMAIN response is sent (depending on the option set for invalid DNS queries, described above). Authoritative domain names are configured with ends-with semantics.

Note: All labels in subdomain and authoritative domain names must be complete. To illustrate by example, suppose alpha.beta.com, delta.beta.com, delta.eta.com, and gamma.eta.com are valid FQDNs. If we intend the GSLB DNS to return authoritative responses to queries for each of the four FQDNs, two authoritative domains could be identified, beta.com and eta.com. It is not sufficient to stipulate eta.com alone because "eta" is not a complete label, and therefore doesn't match either alpha.beta.com or delta.beta.com.

L4 Profile

The L4 Profile is used for any virtual service that does not require application-layer proxying.

Note: Using an L4 profile is equivalent to setting the virtual service's application profile to 'none'.

Rate limits may be placed on the number of TCP connections or UDP packets that may be made to the virtual service from a single client IP address.

- **Threshold ?** The client has violated the rate limit when the defined threshold of connections (TCP) or packets (UDP) is reached within the specified time period.
- **Time Period ?** The client has violated the rate limit when the defined threshold of connections (TCP) or packets (UDP) is reached within the specified time period.
- **Action ?** Select the action to perform when a client has exceeded the rate limit.

- **Report Only ?** A log is generated in the virtual service logs page. By default, no action is taken. However, this option may be used with an alert to generate an alert action to send a notice to a remote destination or to take action using a ControlScript.
- **Drop SYN Packets ?** For TCP-based limits, silently discard TCP SYNs from the client. Avi Vantage also will generate a log. However, during high volumes of DoS traffic, repetitive logs may be skipped.
- **Send TCP RST ?** Reset client TCP connection attempts. While more graceful than the Drop SYN Packet option, sending a TCP reset does generate extra packets for the reset, versus the Drop SYN Packet option which does not send a client response. Avi Vantage also will generate a log. However, during high volumes of DoS traffic, repetitive logs may be skipped.

Syslog Profile

The Syslog application profile allows Avi Vantage to decode the Syslog protocol. This profile will set the virtual service to understanding Syslog, and the network profile to UDP with per-stream parsing.

SIP Profile

SIP profile allows Avi Vantage to process traffic for SIP applications. This profile defines the transaction timeout allowed for SIP traffic through Avi Vantage. Configure the timeout within the range of 16 to 512 seconds.