



HTTP Server Reselect

Avi Technical Reference (v17.2)

Copyright © 2019

HTTP Server Reselect

[view online](#)

Avi Vantage can be configured to retry an HTTP request that fails or returns one of a set of user-specified error codes from the back-end server. Normally, Avi Vantage forwards these error messages back to the client.

Configurable Options

HTTP server reselect is disabled by default. The feature can be configured within individual pools. For releases prior to 17.2.13, when enabling the HTTP server reselect feature, the error codes that trigger the feature need to be specified. Starting with 17.2.13, specification of status codes is not required for the feature to work in connection/SSH failure scenarios.

All other options have defaults that can be changed if needed.

Error Codes

The pool configuration specifies the HTTP error response codes that will result in server reselection. The error codes can be specified in any of the following ways:

- Explicit code number(s): Enter one or more individual error codes (for example, 404).
- Range of codes: Enter a range within 400-499 or 500-599 (for example, 501-503).
- Entire block of codes: 4xx or 5xx (for example, 4xx).

Maximum Retries

The default maximum retry setting is 4. Following the first error response, Avi Vantage will resend the request to the pool up to 4 more times, for a total of 5 attempts. Each retry is sent to a different server, and each server can receive only one attempt.

If the setting for maximum retries is higher than the number of enabled and up servers within the pool, each of those servers still receives only one attempt. For example, if maximum retries is set to 4 but the pool has only 3 servers, the maximum number of retries is actually 2. The initial attempt that fails goes to one of the servers, leaving 2 more servers to try. If the second server also sends a 4xx or 5xx error code in response to the request, the request is sent to the last server in the pool. If the last server also sends a 4xx or 5xx, the response from the server is sent back to the client.

Server Reselection for Idempotent Requests

Currently, HTTP server reselect applies only to idempotent request methods, since a given request of this type always has the same result, even if an identical request is received multiple times by a server. Multiple identical requests of non-idempotent request methods (POST, LOCK, PATCH, and CONNECT) are by definition not guaranteed to have the same effect as a single such request. For this reason, HTTP server reselect is not performed for these request methods.

Configuring HTTP Server Reselect

HTTP server reselect can be enabled on the Advanced tab of the pool configuration.

1. Navigate to Applications > Pools.
2. Open the configuration popup for the pool:
 - If enabling the feature in an existing pool, click the edit icon next to the pool name.
 - If creating a new pool, click Create, then click on the cloud name and click Next. Enter a name for the pool on the Settings tab, then select the servers on the Servers tab.

3. Select Advanced to display the tab of advanced options.
4. Select (check) the HTTP Server Reselect checkbox.
5. Enter the error response codes that will trigger server reselection.
6. Save the changes:
 - If creating a new pool, click Next to review the settings, then click Save.
 - If editing an existing pool, click Save.

The following example enables HTTP server reselection for all 4xx error codes:

The screenshot shows the configuration interface for a pool, specifically the 'Step 3: Advanced' tab. The interface is divided into three main sections: 'Pool A/B Testing Settings', 'Pool Failure Settings', and 'Other Settings'. In the 'Other Settings' section, the 'HTTP Server Reselect' checkbox is checked and highlighted with an orange border. Below this checkbox, the 'HTTP Response codes for Server Reselect' field contains the text '4xx', and the 'Number of Retries' field contains the number '4'. Other visible settings include 'A/B Test Pool' set to 'None', 'Percentage of Requests' set to '0%', 'Pool Fail Action' set to 'Close Connection', 'Disable Port Translation' unchecked, 'Connection Ramp' set to '10 seconds', and 'Max Connections per Server' set to '0'. The bottom of the interface features a 'CANCEL' button, a 'PREVIOUS' button, and a 'NEXT' button.

Based on this configuration, if a server in this pool responds to a client request with a 4xx error code, Avi Vantage retries the request by sending it to another server in the pool. Avi Vantage will retry the request up to 4 times (to 4 different servers).