



Analytics Profile

Avi Technical Reference (v17.1)

Copyright © 2018

Analytics Profile

[view online](#)

Avi Vantage relies extensively on analytics throughout the system to determine the health of applications based on expectations of what a typical user experience should be. Since each application is different, it may be necessary to modify the analytics profile to set the threshold for satisfactory client experience or omit certain errors from being counted against the application health, such as prompting a user to log in to a site via an HTTP 401 response code.

Profile Settings

The following options are available within the analytics profile:

- Name: Enter a unique name for the analytics profile.

- HTTP Analytics: These settings will only be applied to virtual services configured with an HTTP application profile. For non-HTTP virtual services, these HTTP settings will have no effect.

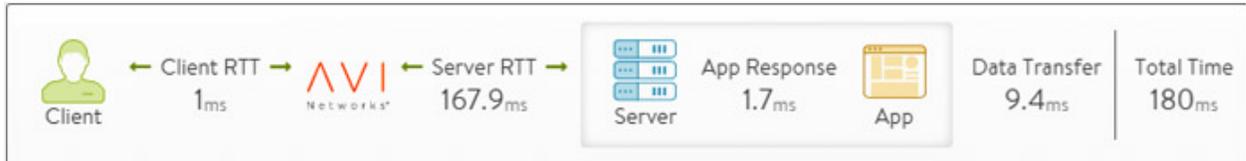
Apdex: The concept of Apdex is used extensively by Avi Vantage to capture a client's experience when accessing a virtual service. Apdex is an industry standard for rating a user's experience, which it classifies as satisfied, tolerated, or frustrated.

0 - 500 ms	501 - 2000 ms	> 2000 ms
Satisfied	Tolerated	Frustrated

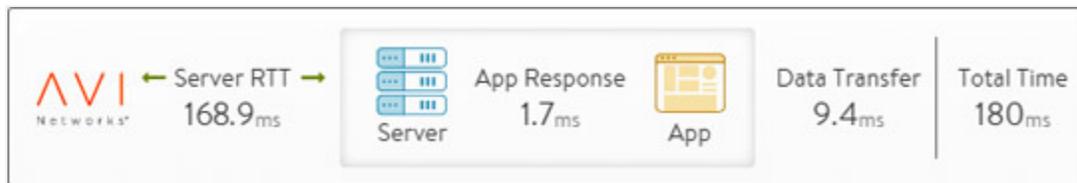
The results are used as part of the performance metric of the pool and virtual service health score. The greater the number of satisfied responses, the higher the score. See www.apdex.org.

- Client Satisfactory Latency Threshold: A client must receive a completed response to an HTTP request within this time frame to be considered satisfied with the transaction time. Using the default of 500ms, any response

that is complete with 0-500ms is considered satisfactory. A response time is measured via the End to End Timing's Total Time metric, which includes Client RTT, Server RTT, App Response, and the Data Transfer metrics.



- **Client Tolerated Latency Factor:** The Satisfied Latency Threshold is multiplied by the Tolerated Factor to determine the tolerated threshold. If satisfied is 500ms and the Tolerated Factor is 4x, then 0-500ms is satisfied, while 500ms to 2000ms is considered a tolerated response time and anything over 2000ms is considered frustrated.
- **Server Satisfactory Latency Threshold:** This is the same as the Client Satisfactory Latency Threshold; however, this metric takes the client's latency out of the picture by only measuring Server RTT, App Response, and the Data Transfer time between the server and the Service Engine. This is similar to viewing the End to End Timing of a pool rather than a virtual service. This metric helps differentiate between poor response times due to slow servers versus slow clients.



- **Server Tolerated Latency Factor:** Similar to the Client Tolerated Latency Factor, this option is multiplies the Server Satisfactory Latency Threshold to determine tolerated responses from servers.
- **Client PageLoad Satisfactory Latency Threshold:** Similar to the Client Satisfactory Latency Threshold, this metric looks at PageLoad times rather than a single HTTP request. PageLoad requires the HTTP virtual service to have the analytics type set to *active*, which will insert JavaScript into a sampling of HTTP responses. PageLoad measures the time it takes for a client to download an entire web page, which may include many objects. It also includes the time for DNS lookups, TCP connection setup, object download, and page rendering. For example, Avi Vantage may see satisfied file transfers for objects it is serving, even though clients are seeing errors due to third-party HTML files being slow or having JavaScript errors. This metric is intended catch these issues and incorporate them into the health score.

This metric will catch these issues.

- **Client PageLoad Tolerated Latency Factor:** Similar to the Client Tolerated Latency Factor, this field is used to determine clients using PageLoad that are having a tolerated or frustrated experience.
- **Exclude HTTP Status Codes from Error Classification:** By default any 4xx and 5xx responses are considered errors. The greater the percentage of errors, the more the performance health score metric is lowered. Errors are also logged via the significant client logs. Some of these errors may need to be excluded for certain applications. For instance, Sharepoint will send a 401 error to clients, asking them to first authenticate before accessing the web site.

- **Network Analytics: Interruptions to TCP connections** may happen for a number of reasons. These connections are deemed lossy and are logged via the significant client logs. They may also reduce the performance health score metric. These metrics are applicable for any virtual service using TCP in proxy mode. Lossy connections may happen on the client or server side of Avi Vantage. Their influence on health scores may be adjusted below.

• Network Analytics •

<p>Client Connection Apdex - Lossy Connection Threshold</p> <p>TCP Retransmit Threshold ?</p> <p>50 % </p> <p>TCP Timeout Threshold ?</p> <p>20 % </p> <p>TCP Out-Of-Order Threshold ?</p> <p>50 % </p> <p>TCP Zero Window Threshold ?</p> <p>1 % </p> <p><input type="checkbox"/> Exclude Client TCP RST as an error ?</p> <p><input type="checkbox"/> Exclude Client Connection Persistence Change as an error ?</p> <p><input type="checkbox"/> Exclude Client Connection close before HTTP Request as an error ?</p>	<p>Server Connection Apdex - Lossy Connection Threshold</p> <p>TCP Retransmit Threshold ?</p> <p>50 % </p> <p>TCP Timeout Threshold ?</p> <p>20 % </p> <p>TCP Out-Of-Order Threshold ?</p> <p>50 % </p> <p>TCP Zero Window Threshold ?</p> <p>1 % </p> <p><input type="checkbox"/> Exclude Server TCP RST as an error ?</p>
--	---

- **TCP Retransmit Threshold:** The TCP connection is considered lossy when more than this percent of packets are retransmitted.
- **TCP Timeout Threshold:** Similar to the previous metric, this option specifically evaluates the number of retransmissions that were required due to timeouts.
- **TCP Out-of-Order Threshold:** The connection is deemed lossy when more than this percentage of packets received from the client are out of order.
- **TCP Zero Window Threshold:** The connection is deemed lossy when greater than this percentage of packets could not be transmitted because the TCP connection window had reduced to zero.
- **Exclude Network Errors:** Some errors may not be abnormal for a given environment. Excluding this issues from the list of errors ensures they will not degrade health score or generate logs.
 - **Client TCP RST:** A graceful TCP shutdown occurs via a FIN/ACK process. Avi Vantage records RST packets as an error unless otherwise excluded.
 - **Client Connection Persistence Change:** Connection persistence change is typically due to a server going offline, forcing Avi Vantage to rebalance connections to new servers. Selecting this option excludes this scenario from the list of errors.
 - **Client Connection Close before HTTP Request:** If the client closes the connection prior to completing an HTTP request, Avi Vantage will record this as an error. Selecting this option excludes this scenario from the list of errors.
 - **Server TCP RST:** A graceful TCP shutdown normally occurs via a FIN/ACK process. Avi Vantage records RST packets as an error unless otherwise excluded. Applications such as Microsoft Exchange may use RST to close connections. Selecting this option will omit server RSTs from the list of errors.

- **Health Score Analytics:** Health scores are assigned to servers, pools, virtual services, and Service Engines. The following settings specifically apply to modifying the health scores of virtual services, whose scores are comprised of performance, anomaly, and resource penalties.

• Health Score Analytics •

<p>Performance Boost ?</p> <div style="border: 1px solid #ccc; padding: 2px; display: flex; align-items: center;"><input style="width: 40px; text-align: center;" type="text" value="0"/>%<div style="border-left: 1px solid #ccc; border-right: 1px solid #ccc; height: 15px; width: 10px; background-color: #007bff; margin: 0 5px;"></div><div style="flex-grow: 1; border-bottom: 1px solid #ccc; position: relative;"><div style="position: absolute; left: -5px; top: -5px; right: -5px; bottom: -5px; background-color: #e0e0e0;"></div></div></div>	<p>Anomaly Penalties ?</p> <div style="border: 1px solid #ccc; padding: 2px; display: flex; align-items: center;"><input style="width: 40px; text-align: center;" type="text" value="10"/>% Maximum</div>	<p>Resource Penalties ?</p> <div style="border: 1px solid #ccc; padding: 2px; display: flex; align-items: center;"><input style="width: 40px; text-align: center;" type="text" value="25"/>% Maximum</div>
--	--	---

- **Performance Boost:** Some applications may simply not be able to consistently meet a response Apdex threshold. For example, an application that relies on a backend database may normally respond within 50ms, but occasional DB queries may take 2 seconds. Rather than set the client and server response Apdex thresholds greater than 2 seconds, instead the performance metric can be artificially inflated by a small percentage. This allows the satisfied threshold to remain aggressive while still allowing for occasional slow responses.
- **Anomaly Penalties:** This setting controls how many points Anomalies will deduct from the performance score. Anomalies represent risk to the application via inconsistent behavior of clients, traffic volume, or server responses. Lowering the anomaly penalty places greater emphasis on the performance score and resource penalties.
- **Resource Penalties:** Resources that are constrained will increase the resource penalty score. Examples include the CPU, memory, or disk utilization of a Service Engine or a server running on a virtual machine.