



Metrics List

Avi Technical Reference (v17.2)

Copyright © 2020

Metrics List

[view online](#)

Avi Networks continuously adds additional metrics to Avi Vantage. They are used to trigger alerts, notify administrators of important information, and are the building blocks of much of the configuration automation. See [Events Overview](#), [Alerts Overview](#), and [Notifications Overview](#) for a larger view of how the event component is used for alerting and automation.

Object Types

At the end of each metric's description below, the types of objects to which the metric applies are listed. Metrics apply to one or more of the following types of objects:

- VS: Metric applies to virtual services.
- SE: Metric applies to Service Engines.
- P: Metric applies to pools.

Layer 4 Metrics

- `I4_client.apdexc`: Measures the network connection quality (errors and lossy connections) between clients and the VS. (VS)
- `I4_client.apdexrtt`: Measures the network connection quality against the Round Trip Time between clients and the VS. (VS)
- `I4_client.avg_application_dos_attacks`: Number of application layer DoS attacks occurring. (VS)
- `I4_client.avg_bandwidth`: Average transmit and receive network bandwidth between client and virtual service. (VS)
- `I4_client.avg_complete_conns`: Rate of total connections completed (closed) per second. (VS)
- `I4_client.avg_connections_dropped`: Rate of dropped connections per second. (VS)
- `I4_client.avg_dos_attacks`: Number of application plus network layer DoS attacks occurring. (VS)
- `I4_client.avg_dos_bandwidth`: Average transmit and receive network bandwidth between client and virtual service related to DoS attack. (VS)
- `I4_client.avg_errored_connections`: Rate of connections per second classified as errored, having network connection errors. (VS)
- `I4_client.avg_I4_client_latency`: Average L4 connection duration which does not include client RTT. (VS)
- `I4_client.avg_lossy_connections`: Rate of connections per second classified as lossy, having significant packet loss and retransmits. (VS)
- `I4_client.avg_network_dos_attacks`: Number of network layer DoS attacks occurring. (VS)
- `I4_client.avg_new_established_conns`: New client connections per second. (VS)
- `I4_client.avg_policy_drops`: Rate of connections dropped per second due to virtual service policies. This includes rate limits, security policies, and connection limits. (VS)
- `I4_client.avg_total_rtt`: Average network round trip time between client and virtual service. (VS)
- `I4_client.max_open_conns`: Maximum number of concurrently open connections. (VS)
- `I4_client.pct_application_dos_attacks`: Percentage of HTTP application requests classified as a DoS attack. (VS)
- `I4_client.pct_connection_errors`: Percentage of network connections to the virtual service classified as errored due to dropped or lossy. (VS)
- `I4_client.pct_connections_dos_attacks`: Percentage of network connections classified as a DoS attack. (VS)
- `I4_client.pct_dos_bandwidth`: Percentage of network bandwidth classified as part of a DoS attack. (VS)
- `I4_client.pct_dos_rx_bytes`: Percentage of bytes received that are classified as part of a DoS attack. (VS)
- `I4_client.pct_pkts_dos_attacks`: Percentage of packets that are classified as part of a DoS attack. (VS)
- `I4_client.pct_policy_drops`: Percentage of network connections dropped due to a virtual service network security policy. (VS)
- `I4_client.sum_connection_errors`: Total number of client network connections that were lossy or dropped. (VS)

- `I4_client.sum_finished_conns`: Total number of completed connections. (VS)
- `I4_client.sum_lossy_connections`: Total connections that were classified as lossy due to high packet retransmissions. (VS)
- `I4_client.sum_lossy_req`: Total HTTP requests that were classified as lossy due to high packet retransmissions. (VS)
- `I4_server.apdexc`: Measures the network connection quality (errors and lossy connections) between Service Engines and servers. (VS, P)
- `I4_server.avg_available_capacity`: Estimated connections per second capacity available for a server. This metric is the difference between the max capacity and the averaged load on a server. (VS, P)
- `I4_server.avg_bandwidth`: Average transmit and receive network bandwidth between client and virtual service. (VS, P)
- `I4_server.avg_complete_conns`: Rate of new connections per second. (VS, P)
- `I4_server.avg_connections_dropped`: Rate of dropped connections per second. (VS, P)
- `I4_server.avg_errored_connections`: Rate of total errored connections per second. (VS, P)
- `I4_server.avg_est_capacity`: Estimated averaged capacity of a server's connections per second summed across all Service Engines. Pool level metric reflects summed estimated capacity across all the servers in the pool. (VS, P)
- `I4_server.avg_goodput`: Application data goodput (data excluding network headers) as bytes per second between the Service Engine and server. (VS, P)
- `I4_server.avg_health_status`: Health score status of the server. 0 is down, 1 is up though with significant issues, 100 is ideal. (VS, P)
- `I4_server.avg_lossy_connections`: Rate of lossy connections per second between the Service Engine and server. (VS, P)
- `I4_server.avg_new_established_conns`: Rate of new established connections per second between Service Engine and server. (VS, P)
- `I4_server.avg_open_conns`: Number of concurrently open connections between Service Engines and servers. (VS, P)
- `I4_server.avg_pool_bandwidth`: Transmit and receive network bandwidth between Service Engines and all servers in a pool. (VS, P)
- `I4_server.avg_pool_complete_conns`: New connections per second across the VS or pool, divided by the number of active (up state) servers. (VS, P)
- `I4_server.avg_pool_errored_connections`: Total connections classified as errored between Service Engines and all servers in a pool. (VS, P)
- `I4_server.avg_pool_new_established_conns`: Total new connections per second established between Service Engines and all servers in a pool. (VS, P)
- `I4_server.avg_total_rtt`: Average Round Trip Time across all completed (closed) connections. (VS, P)
- `I4_server.avg_uptime`: Percent of time a server was marked as up. (VS, P)
- `I4_server.max_open_conns`: Maximum number of concurrently open connections to a server. (VS, P)
- `I4_server.pct_connection_errors`: Percent of network connections between Service Engines and a server that were dropped or lossy. (VS, P)
- `I4_server.pct_connection_saturation`: Percent of a server's connection per second capacity that is estimated to be utilized. (VS, P)
- `I4_server.sum_connection_errors`: Total number of network connections to a server that were dropped or were classified as lossy. (VS, P)
- `I4_server.sum_connections_dropped`: Total number of network connections to a server that were dropped. (VS, P)
- `I4_server.sum_finished_conns`: Total number of completed connections to a server. (VS, P)
- `I4_server.sum_health_check_failures`: Total number of times a server was marked down by health monitors. (VS, P)
- `I4_server.sum_lossy_connections`: Total number of network connections to a server that were classified as lossy. (VS, P)
- `I4_server.sum_lossy_req`: Total number of HTTP requests that were classified as lossy due to high packet retransmissions. (VS, P)

Layer 7 Metrics

- `I7_client.apdexr`: Measures the clients view of the quality (response time and errors) of server responses. (VS, P)
- `I7_client.avg_application_response_time`: Average server / application response latency. (VS, P)

- `l7_client.avg_blocking_time`: Average time client was blocked as reported by client. Requires Client Insights set to Active. (VS)
- `l7_client.avg_browser_rendering_time`: Average time browser spend rendering a web page. Requires Client Insights set to Active. (VS)
- `l7_client.avg_client_data_transfer_time`: Average time required to transmit a file via layer 7 protocol (such as HTTP) from a virtual service to a client, excluding the Round Trip Time. (VS)
- `l7_client.avg_client_rtt`: Average Round Trip Time between client and virtual service. (VS)
- `l7_client.avg_client_txn_latency`: Client transaction time averaged across all HTTP requests. (VS)
- `l7_client.avg_complete_responses`: Average rate of HTTP responses sent to clients per second. (VS)
- `l7_client.avg_connection_time`: Average client connection latency as reported by client. Requires Client Insights set to Active to gather RUM data. (VS)
- `l7_client.avg_dns_lookup_time`: Average DNS name lookup time as reported by the client. Requires Client Insights set to Active to gather RUM data. (VS)
- `l7_client.avg_dom_content_load_time`: Average Dom content load time as reported by clients. Requires Client Insights set to Active to gather RUM data. (VS)
- `l7_client.avg_error_responses`: Rate of HTTP error responses per second sent to clients. It does not include error codes that have been excluded in analytics profile. (VS)
- `l7_client.avg_frustrated_responses`: Number of client HTTP requests that are completed but classified in the Frustrated latency bucket. (VS)
- `l7_client.avg_http_headers_bytes`: Average size of HTTP headers per request.
- `l7_client.avg_http_headers_count`: Average number of HTTP headers per request.
- `l7_client.avg_http_params_count`: Average number of HTTP request parameters per request.
- `l7_client.avg_page_download_time`: Page download time as reported by clients. Requires Client Insights set to Active. (VS)
- `l7_client.avg_page_load_time`: Page load time as reported by clients. Requires Client Insights set to Active. (VS)
- `l7_client.avg_params_per_req`: Average number of HTTP request parameters per request, taking into account only requests with parameters.
- `l7_client.avg_post_bytes`: Average size of HTTP POST request.
- `l7_client.avg_redirection_time`: Latency incurred by following redirects as reported by clients. Requires Client Insights set to Active. (VS)
- `l7_client.avg_resp_1xx`: Rate of HTTP 1xx responses sent to clients. (VS)
- `l7_client.avg_resp_2xx`: Rate of HTTP 2xx responses sent to clients. (VS)
- `l7_client.avg_resp_3xx`: Rate of HTTP 3xx responses sent to clients. (VS)
- `l7_client.avg_resp_4xx`: Rate of HTTP 4xx responses sent to clients. (VS)
- `l7_client.avg_resp_4xx_avi_errors`: Rate of HTTP 4xx responses sent to clients from Avi, such as via custom security policies. This does not include errors excluded via the analytics profile or server generated 4xx errors. (VS)
- `l7_client.avg_resp_5xx`: Rate of HTTP 5xx responses sent to clients. (VS)
- `l7_client.avg_resp_5xx_avi_errors`: Rate of HTTP 5xx responses sent to clients from Avi, such as via custom security policies. This does not include errors excluded via the analytics profile or server generated 5xx errors. (VS)
- `l7_client.avg_rum_client_data_transfer_time`: Total client data transfer time as reported by clients. Requires Client Insights set to Active to gather RUM data. (VS)
- `l7_client.avg_satisfactory_responses`: Number of client HTTP requests that are completed and classified in the Satisfied latency bucket. (VS)
- `l7_client.avg_server_rtt`: Average Round Trip Time between the Service Engine and server. (VS)
- `l7_client.avg_service_time`: Average latency from the virtual services receipt of a request to start of the response. Requires Client Insights set to Active to gather RUM data. (VS)
- `l7_client.avg_ssl_connections`: New SSL transactions per second (TPS), including SSL session reuse and failed handshake negotiations. (VS)
- `l7_client.avg_ssl_errors`: SSL connection errors per second due to clients, protocol errors, network errors and handshake timeouts (VS)

- `I7_client.avg_ssl_failed_connections`: SSL connection errors per second due to protocol, network, or timeout issues. (VS)
- `I7_client.avg_ssl_handshakes_new`: New SSL transactions per second (TPS), excluding session reuse and errored connection attempts. (VS)
- `I7_client.avg_ssl_handshakes_non_pfs`: New SSL handshakes / transactions per second (TPS) that did not use Perfect Forward Secrecy. (VS)
- `I7_client.avg_ssl_handshakes_pfs`: New SSL handshakes / transactions per second (TPS) that used Perfect Forward Secrecy. (VS)
- `I7_client.avg_ssl_handshakes_reused`: Successfully resumed SSL sessions per second. (VS)
- `I7_client.avg_tolerated_responses`: Number of HTTP requests which had response latency classified as Tolerated per the virtual service analytics profile. (VS)
- `I7_client.avg_total_requests`: Client HTTP requests per second received by the virtual service. (VS)
- `I7_client.avg_uri_length`: Average length of HTTP URI per request.
- `I7_client.avg_waf_disabled`: Average number of transactions per second bypassing WAF.
- `I7_client.avg_waiting_time`: Average waiting time reported by the client. Requires Client Insights set to Active. (VS)
- `I7_client.max_ssl_open_sessions`: Maximum number of concurrently open SSL sessions. (VS)
- `I7_client.pct_cache_hits`: Percent of HTTP requests served from cache. (VS)
- `I7_client.pct_cacheable_hits`: Percent of HTTP requests that were eligible to be served from cache. (VS)
- `I7_client.pct_get_reqs`: Number of HTTP GET requests as a percentage of total requests received.
- `I7_client.pct_post_reqs`: Number of HTTP POST requests as a percentage of total requests received.
- `I7_client.pct_response_errors`: Percent of 4xx and 5xx HTTP responses. (VS)
- `I7_client.pct_ssl_failed_connections`: Percent of SSL connection failures due to protocol, network, or timeout errors. (VS)
- `I7_client.pct_waf_disabled`: Transactions bypassing WAF as the percentage of total requests received.
- `I7_client.rum_apdexr`: Quality (combination of performance and errors) of HTTP responses to clients based on RUM data. Requires Client Insights set to Active for RUM data. (VS)
- `I7_client.sum_errors`: Total number of HTTP 400 and 500 errors sent to a client. (VS)
- `I7_client.sum_get_reqs`: Total number of HTTP GET requests. (VS)
- `I7_client.sum_http_headers_bytes`: Total size of HTTP request headers in a given metrics interval.
- `I7_client.sum_http_headers_count`: Total number of HTTP headers across all requests in a given metrics interval.
- `I7_client.sum_http_params_count`: Total number of HTTP request parameters.
- `I7_client.sum_num_rum_samples`: Total number of samples used for RUM metrics. Requires Client Insights set to Active to gather RUM data. (VS)
- `I7_client.sum_other_reqs`: Total number of HTTP requests that are not GET or POST requests. (VS)
- `I7_client.sum_post_bytes`: Total size of HTTP POST requests.
- `I7_client.sum_post_reqs`: Total number of HTTP POST requests. (VS)
- `I7_client.sum_reqs_with_params`: Total number of HTTP requests containing at least one parameter.
- `I7_client.sum_total_responses`: Total number of HTTP responses sent to clients. (VS)
- `I7_client.sum_uri_length`: Total length of HTTP request URIs.
- `I7_client.sum_waf_disabled`: Total number of requests bypassing WAF in a given metrics interval.
- `I7_server.apdexr`: Quality (combination of performance and errors) of HTTP responses from servers to the virtual service. (VS)
- `I7_server.avg_application_response_time`: Average response latency measured of pool servers. (VS)
- `I7_server.avg_complete_responses`: Rate of server HTTP responses per second. (VS, P)
- `I7_server.avg_error_responses`: Rate of HTTP error responses sent per second. Does not include errors excluded in analytics profile. (VS, P)
- `I7_server.avg_frustrated_responses`: Number of HTTP requests completed which had server response latency classified as Frustrated per the virtual service analytics profile. (VS, P)
- `I7_server.avg_resp_1xx`: Rate of 1xx HTTP responses sent per second. (VS, P)
- `I7_server.avg_resp_2xx`: Rate of 2xx HTTP responses sent per second. (VS, P)
- `I7_server.avg_resp_3xx`: Rate of 3xx HTTP responses sent per second. (VS, P)

- `l7_server.avg_resp_4xx`: Rate of 4xx HTTP responses sent per second. (VS, P)
- `l7_server.avg_resp_4xx_errors`: Rate of 4xx HTTP responses per second minus error codes excluded by the analytics profile. (VS, P)
- `l7_server.avg_resp_5xx`: Rate of 5xx HTTP responses sent per second. (VS, P)
- `l7_server.avg_resp_5xx_errors`: Rate of 5xx HTTP responses per second minus error codes excluded by the analytics profile. (VS, P)
- `l7_server.avg_resp_latency`: Latency measured for pool servers. (VS, P)
- `l7_server.avg_satisfactory_responses`: Number of HTTP requests completed which had server response latency classified as Satisfied per the virtual service analytics profile. (VS, P)
- `l7_server.avg_tolerated_responses`: Number of HTTP requests completed which had server response latency classified as Tolerated per the virtual service analytics profile. (VS, P)
- `l7_server.avg_total_requests`: Rate of HTTP requests per second received by pool servers. (VS, P)
- `l7_server.pct_response_errors`: Percent of HTTP 4xx and 5xx server responses. (VS, P)
- `l7_server.sum_get_reqs`: Total number of HTTP GET requests received by servers. (VS, P)
- `l7_server.sum_other_reqs`: Total number of HTTP requests that are not GET or POST request received by servers. (VS, P)
- `l7_server.sum_post_reqs`: Total number of HTTP POST requests received by servers. (VS, P)
- `l7_server.sum_total_responses`: Total number of HTTP responses sent from servers. (VS, P)

Service Engine Metrics

- `se_if.avg_bandwidth`: Transmit and receive network bandwidth across all Service Engine interfaces. (VS, SE)
- `se_stats.avg_connection_mem_usage`: Percent of connection memory consumed. (VS, SE)
- `se_stats.avg_connections`: Rate of client network connection attempts (SYNs) per second across all virtual services for a Service Engine. (VS, SE)
- `se_stats.avg_connections_dropped`: Number of connections dropped or failed to establish across all virtual services for a Service Engine. Excludes drops due to a policy action. (VS, SE)
- `se_stats.avg_cpu_usage`: Hosts view of the Service Engines actively utilized CPU usage as a percent of total available CPU. (VS, SE)
- `se_stats.avg_disk1_usage`: Percent of Service Engine disk 1 capacity used (VS, SE)
- `se_if.avg_eth0_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth10_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth11_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth12_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth13_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth14_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth15_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth16_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth17_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth18_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth19_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth1_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth20_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth21_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth22_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth23_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth2_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth3_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth4_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth5_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth6_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)

- `se_if.avg_eth7_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth8_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_if.avg_eth9_bandwidth`: Transmit and receive network bandwidth for a Service Engine interface. (VS, SE)
- `se_stats.avg_mem_usage`: Percent of allocated memory used (VS, SE)
- `se_stats.avg_packet_buffer_header_usage`: Percent of all network packet buffers used. (VS, SE)
- `se_stats.avg_packet_buffer_large_usage`: Percent of large network packet buffers used. (VS, SE)
- `se_stats.avg_packet_buffer_small_usage`: Percent of small network packet buffers used. (VS, SE)
- `se_stats.avg_packet_buffer_usage`: Percent of total configured network packet buffers used. (VS, SE)
- `se_stats.avg_persistent_table_usage`: Percent of session persistence table used. (VS, SE)
- `se_if.avg_rx_bandwidth`: Received network bandwidth across all interfaces for a Service Engine. (VS, SE)
- `se_if.avg_rx_bytes_dropped`: Bytes per second of received packets that were dropped. Includes packets across all Service Engine interfaces. (VS, SE)
- `se_if.avg_rx_pkts`: Received network packets across all Service Engine interfaces. (VS, SE)
- `se_if.avg_rx_pkts_dropped`: Received packets dropped per second. Includes packet drops across all interfaces. (VS, SE)
- `se_stats.avg_ssl_session_cache_usage`: Percent of SSL session cache used. (VS, SE)
- `se_if.avg_tx_pkts`: Network packets transmitted across all Service Engine interfaces. (VS, SE)
- `se_if.max_se_bandwidth`: Maximum bandwidth through a Service Engine within the sample period. (VS, SE)
- `se_stats.pct_connections_dropped`: Percent of Service Engine connections dropped. (VS, SE)
- `se_if.pct_rx_bytes_dropped`: Percent of Service Engine bytes dropped. (VS, SE)
- `se_if.pct_rx_pkts_dropped`: Percent of Service Engine packets dropped. (VS, SE)
- `se_stats.pct_syn_cache_usage`: Percent of SYN cache used. Higher usage indicates too many incomplete open connection attempts. (VS, SE)

Insight Metrics

- `source_insights.avg_bandwidth`: Rate of network throughput. (VS)
- `source_insights.avg_client_end2end_latency`: Client end to end latency for successful responses. (VS)
- `source_insights.avg_complete_conns`: Rate of connections per second. (VS)
- `source_insights.avg_complete_responses`: Rate of completed responses including response errors. (VS)
- `source_insights.avg_dropped_conns`: Rate of dropped connections per second. (VS)
- `source_insights.avg_error_responses`: Rate of errored requests per second. (VS)
- `source_insights.avg_http_timeout`: Rate of HTTP timeouts per second. (VS)
- `source_insights.avg_policy_drops`: Rate of connections per second dropped by policies. (VS)
- `source_insights.avg_total_requests`: Rate of HTTP requests per second. (VS)
- `source_insights.max_open_conns`: Maximum number of concurrent open connections. (VS)

VMware Metrics

- `vm_stats.avg_cpu_usage`: Percent of server CPU used. (VS, SE)
- `vm_stats.avg_cpu_wait`: Percent of time virtual machine was ready to run but could not due to CPU unavailable. This could be CPU limits configured in vCenter or other virtual machines stealing CPU time. (VS, SE)
- `vm_stats.avg_disk1_usage`: Virtual disk1 capacity usage. (VS, SE)
- `vm_stats.avg_disk2_usage`: Virtual disk2 capacity usage. (VS, SE)
- `vm_stats.avg_disk3_usage`: Virtual disk3 capacity usage. (VS, SE)
- `vm_stats.avg_disk4_usage`: Virtual disk4 capacity usage. (VS, SE)
- `vm_stats.avg_disk_commands_aborted`: Rate of disk I/O commands that were aborted in a virtual machine. (VS, SE)
- `vm_stats.avg_disk_io`: Rate of server disk reads plus writes per second. (VS, SE)
- `vm_stats.avg_disk_read`: Rate of data read from disk in kilobytes per second. (VS, SE)
- `vm_stats.avg_disk_write`: Rate of data written to disk in kilobytes per second. (VS, SE)
- `vm_stats.avg_mem_swap_in`: Total amount of data that has been read into machine memory from the swap file since the virtual machine was powered on. (VS, SE)

- `vm_stats.avg_mem_swap_out`: Total amount of data the VMkernel has written to the virtual machine swap file from machine memory. This statistic refers to VMkernel swapping and not to guest OS swapping. (VS, SE)
- `vm_stats.avg_mem_usage`: Percent of available server memory used. (VS, SE)
- `vm_stats.avg_mem_vmmemctl`: Virtual machine guests physical memory currently reclaimed from the virtual machine through ballooning. This is the amount of guest physical memory that has been allocated and pinned by the balloon driver. (VS, SE)
- `vm_stats.avg_net_dropped`: Rate of dropped received and transmit packets. May be an indication of network congestion. (VS, SE)
- `vm_stats.avg_net_usage`: Transmit plus receive network bandwidth for the virtual machine. (VS, SE)
- `vm_stats.avg_port_usage`: Percent of high ports used. (VS, SE)
- `vm_stats.avg_uptime`: Percent of time virtual machine was up during a time interval. (VS, SE)
- `vm_stats.avg_virtual_disk_commands_aborted`: Number of SCSI commands that were aborted. (VS, SE)

Essential Metrics to Monitor

If you have just started to deploy Avi Vantage, the following list helps you to implement an essential list of metrics to monitor:

- VS L4 Client Stats
 - `l4_client.apdexc`
 - `l4_client.avg_bandwidth`
 - `l4_client.avg_complete_conns`
 - `l4_client.max_open_conns`
 - `l4_client.pct_connection_errors`
 - `l4_client.pct_policy_drops`
 - `l4_client.avg_new_established_conns`
 - `l4_client.avg_total_rtt`
- VS L7 Client Stats
 - `l7_client.apdexr`
 - `l7_client.avg_client_txn_latency`
 - `l7_client.avg_complete_responses`
 - `l7_client.avg_error_responses`
 - `l7_client.avg_ssl_connections`
 - `l7_client.avg_ssl_errors`
 - `l7_client.avg_waf_attacks`
 - `l7_client.pct_cache_hits`
 - `l7_client.pct_response_errors`
 - `l7_client.pct_ssl_failed_connections`
- VS L4 Aggregated Server Stats
 - `l4_server.apdexc`
 - `l4_server.avg_bandwidth`
 - `l4_server.avg_complete_conns`
 - `l4_server.max_open_conns`
 - `l4_server.pct_connection_errors`
- VS L7 Aggregated Server Stats
 - `l7_server.apdexr`
 - `l7_server.avg_resp_latency`
 - `l7_client.avg_complete_responses`
 - `l7_server.avg_error_responses`
 - `l7_server.pct_response_errors`
- SE Stats
 - `se_stats.avg_connection_mem_usage`
 - `se_stats.avg_cpu_usage`

- se_stats.avg_mem_usage
- se_stats.avg_packet_buffer_header_usage
- se_stats.avg_packet_buffer_large_usage
- se_stats.avg_packet_buffer_small_usage
- se_stats.avg_packet_buffer_usage
- se_stats.avg_persistent_table_usage
- se_stats.avg_ssl_session_cache_usage
- se_stats.pct_syn_cache_usage
- se_stats.avg_dynamic_mem_usage