



What is Avi's new on-demand approach to storing and processing significant logs?

Avi Technical Reference (v18.1)

What is Avi's new on-demand approach to storing and processing significant logs?

[view online](#)

Answer

By default, significant logs are collected by the Avi Controller *only* when a user sends a log query to the Avi Controller, causing it to fetch the needed log data from the appropriate SEs, index the data, and respond to the user's request. Note that:

- * Absent any queries, logs go no further than the SEs that have collected them. Oldest log entries are rotated out (and their disk space freed) on a per-SE basis when/if an SE runs out of disk space.
- * Logs gathered by a Controller are replicated across all Controllers within the cluster.
- * Prior to an SE being removed, Avi uploads its log files to the Avi Controller.
- * A new option (`logs_processing_auto_sync_and_index`) in the analytics profile is accessible via the CLI. Use it to retain Avi Vantage's original approach, whereby the significant logs captured and stored by SEs are regularly uploaded to the Controller and indexed, and the corresponding SE disk space freed. This option is illustrated by this Avi CLI sequence: `{%cli%}: > configure analyticsprofile mycustomprofile : analyticsprofile> client_log_config : analyticsprofile:client_log_config> significant_log_processing logs_processing_auto_sync_and_index : analyticsprofile:client_log_config> save : analyticsprofile> save {%endcli%}`

The below excerpt shows significant-log collection turned on and processed in the original fashion: `{%cli%}`

```

|-----+-----| enable_significant_log_collection | True | |
significant_log_processing | LOGS_PROCESSING_AUTO_SYNC_AND_INDEX | | filtered_log_processing |
LOGS_PROCESSING_SYNC_AND_INDEX_ON_DEMAND | | non_significant_log_processing |
LOGS_PROCESSING_SYNC_AND_INDEX_ON_DEMAND | | exclude_server_dns_error_as_error | False |
|-----+-----| {%endcli%}

```