

Configuration

Metrics Collector

Configuration Type	File path	Comment
ams-site	/etc/ambari-metrics-collector/conf/ams-site.xml	Settings that control the API daemon and the aggregator threads.
ams-env	/etc/ambari-metrics-collector/conf/ams-env.sh	Memory / PATH settings for the API daemon
ams-hbase-site	/etc/ams-hbase/conf/hbase-site.xml /etc/ambari-metrics-collector/conf/hbase-site.xml	Settings for the HBase storage used for the metrics data.
ams-hbase-env	/etc/ams-hbase/conf/hbase-env.sh	Memory / PATH settings for the HBase storage. Note: In embedded more, the heap memory setting for master and regionserver is summed up as total memory for single HBase daemon.

Metrics Monitor

Configuration Type	File path	Comment
ams-env	/etc/ambari-metrics-monitor/conf/ams-env.sh	Used for log and pid dir modifications, this is the same configuration as above, common to both components.
metric_groups	/etc/ambari-metrics-monitor/conf/metric_groups.conf	Not available in the UI. Used to control what HOST/SYSTEM metrics are reported.
metric_monitor	/etc/ambari-metrics-monitor/conf/metric_monitor.ini	Not available in the UI. Settings for the monitor daemon.

Metric Collector - *ams-site* - Configuration details

- Modifying retention interval for time aggregated data. Refer to Aggregation section for more information on aggregation: [API spec](#) (**Note:** In Ambari 2.0 and 2.1, the Phoenix version does not support Alter TTL queries. So these can be modified from the UI, only at install time. Please refer to [Known Issues](#) section for workaround.)

Property	Default value	Description
timeline.metrics.host.aggregator.ttl	86400	1 minute resolution data purge interval. Default is 1 day.
timeline.metrics.host.aggregator.minute.ttl	604800	Host based X minutes resolution data purge interval. Default is 7 days. (X = configurable interval, default interval is 2 minutes)
timeline.metrics.host.aggregator.hourly.ttl	2592000	Host based hourly resolution data purge interval. Default is 30 days.
timeline.metrics.host.aggregator.daily.ttl	31536000	Host based daily resolution data purge interval. Default is 1 year.
timeline.metrics.cluster.aggregator.minute.ttl	2592000	Cluster wide minute resolution data purge interval. Default is 30 days.
timeline.metrics.cluster.aggregator.hourly.ttl	31536000	Cluster wide hourly resolution data purge interval. Default is 1 year.
timeline.metrics.cluster.aggregator.daily.ttl	63072000	Cluster wide daily resolution data purge interval. Default is 2 years.

Note: The precision table at 1 minute resolution stores raw precision data for 1 day, when user queries for past 1 hour of data, the AMS API returns raw precision data.

- Modifying the aggregation intervals for HOST and CLUSTER aggregators.

On wake up the aggregator threads resume from (last run time + interval) as long as last run time is not too old.

Property	Default value	Description
timeline.metrics.host.aggregator.minute.interval	120	Time in seconds to sleep for the minute resolution host based aggregator. Default resolution is 2 minutes.
timeline.metrics.host.aggregator.hourly.interval	3600	Time in seconds to sleep for the hourly resolution host based aggregator. Default resolution is 1 hour.
timeline.metrics.host.aggregator.daily.interval	86400	Time in seconds to sleep for the day resolution host based aggregator. Default resolution is 24 hours.
timeline.metrics.cluster.aggregator.minute.interval	120	Time in seconds to sleep for the minute resolution cluster wide aggregator. Default resolution is 2 minutes.
timeline.metrics.cluster.aggregator.hourly.interval	3600	Time in seconds to sleep for the hourly resolution cluster wide aggregator. Default is 1 hour.
timeline.metrics.cluster.aggregator.daily.interval	86400	Time in seconds to sleep for the day resolution cluster wide aggregator. Default is 24 hours.

- Modifying checkpoint information. The aggregators store the timestamp or last run time on local FS. After reading last run time, the aggregator thread decides to aggregate as long as the $(currentTime - lastRunTime) < multiplier * aggregation_interval$. The multiplier is configurable for each aggregator.

Property	Default value	Description
timeline.metrics.host.aggregator.minute.checkpointCutOffMultiplier	2	Multiplier value * interval = Max allowed checkpoint lag. Effectively if aggregator checkpoint is greater than max allowed checkpoint delay, the checkpoint will be discarded by the aggregator.
timeline.metrics.host.aggregator.hourly.checkpointCutOffMultiplier	2	Same as above
timeline.metrics.host.aggregator.daily.checkpointCutOffMultiplier	1	Same as above
timeline.metrics.cluster.aggregator.minute.checkpointCutOffMultiplier	2	Same as above
timeline.metrics.cluster.aggregator.hourly.checkpointCutOffMultiplier	2	Same as above
timeline.metrics.cluster.aggregator.daily.checkpointCutOffMultiplier	1	Same as above
timeline.metrics.aggregator.checkpoint.dir	/var/lib/ambari-metrics-collector/checkpoint	Directory to store aggregator checkpoints. Change to a permanent location so that checkpoint are not lost.

- Other important configuration properties

Property	Default value	Description
timeline.metrics.host.aggregator.*.disabled	false	Disable host based * aggregations. (* => minute/hourly/daily)

timeline.metrics.cluster.aggregator.*.disabled	false	Disable cluster based * aggregations. (* => minute/hourly/daily)
timeline.metrics.cluster.aggregator.minute.timeslice.interval	30	Lowest resolution of desired data for cluster level minute aggregates.
timeline.metrics.hbase.data.block.encoding	FAST_DIFF	Codecs are enabled on a table by setting the DATA_BLOCK_ENCODING property. Default encoding is FAST_DIFF. This can be changed only before creating tables.
timeline.metrics.hbase.compression.scheme	SNAPPY	Compression codes need to be installed and available before setting the scheme. Default compression is SNAPPY. Disable by setting to None. This can be changed only before creating tables.
timeline.metrics.service.default.result.limit	5760	Max result limit on number of rows returned. Calculated as follows: 4 aggregate metrics/min * 60 * 24: Retrieve aggregate data for 1 day.
timeline.metrics.service.checkpointDelay	60	Time in seconds to sleep on the first run or when the checkpoint is too old.
timeline.metrics.service.resultset.fetchSize	2000	JDBC resultset prefetch size for aggregator queries.
timeline.metrics.service.cluster.aggregator.applds	datanode, nodemanager,hbase	List of application ids to use for aggregating host level metrics for an application. Example: bytes_read across Yarn Nodemanagers.