

# Known Issues

| No. | JIRA                         | Affected Version (s) | Fix Version | Issue   | Steps to workaround this issue   |
|-----|------------------------------|----------------------|-------------|---|--|
| 14  | -                            | 2.4.x +              | -           | <p>Hiveserver2 can send a ton of metrics which causes performance issues with Metrics Collector.</p> <p><b>Symptoms</b></p> <ul style="list-style-type: none"> <li>AMS goes down intermittently.</li> <li>AMS displays only host metrics (Host summary page on Ambari / System - Servers Grafana dashboard).</li> <li>Aggregated data is not seen. (AMS Summary page / System - Home Grafana dashboard / HBase - Home Grafana dashboard).</li> </ul> <p><b>How do you find out if this is the issue?</b></p> <ul style="list-style-type: none"> <li>Check out AMS metadata endpoint - <code>http://&lt;ams-host&gt;:6188/ws/v1/timeline /metrics/metadata</code>.</li> <li>Use a JSON viewer to look at the JSON response. The number of metrics coming from every component is seen. AMS can handle around ~10000 unique metrics. Check out hiveserver2 or any other component's metric count is causing an explosion of metrics.</li> </ul> | <ul style="list-style-type: none"> <li>Try increasing the heap settings for Metrics collector and HBase (<a href="#">Configurations - Tuning</a>), since the system might be tuned for a small cluster while in fact is getting a lot of metrics.</li> <li>Set <ul style="list-style-type: none"> <li>ams-site : timeline.metrics.service.resultset.fetchSize = 10000</li> <li>ams-hbase-site : ams-hbase-site: hbase.regionserver.handler.count=30</li> </ul> </li> <li>If Ambari 2.5.x, set timeline.metrics.cluster.aggregation.sql.filters = sdisk_%,boottime,default,General%</li> <li>Whitelist - <ul style="list-style-type: none"> <li>In Ambari 2.4.x the whitelisting feature is quite limited (only a metric whitelist file can be used).</li> <li>However in Ambari 2.5.0, a lot of refinements were added to whitelisting.</li> <li>Check out <a href="#">Ambari Metrics - Whitelisting</a>.</li> </ul> </li> </ul> |
| 13  | <a href="#">AMBARI-20056</a> | 2.2.2                | 2.5.0       | <p>On large clusters AMS can become in-operational due to store file explosion and no compaction.</p> <p>Consequence : Large number of store files (~10000) in AMS HBase and AMS shutting down regularly.</p>   | <ol style="list-style-type: none"> <li>Set <code>ams-site : timeline.metrics.hbase.compaction.enabled = false</code></li> <li>Connect to HBase shell on Metrics collector host.<br/><code>/usr/lib/ams-hbase/bin/hbase -config /etc/ams-hbase/conf shell</code></li> <li>Execute the following statement</li> </ol> <pre>alter 'METRIC_RECORD', CONFIGURATION =&gt; { 'hbase.hstore.blockingStoreFiles' =&gt; '1000', 'hbase.hstore.defaultengine.compactionpolicy.class' =&gt; 'org.apache.hadoop.hbase.regionserver.Compactions.FIFOCompactionPolicy' }</pre> <p>If the above does not solve the issue, the only way to recover the system is to reset the metric system.</p>  |
| 12  | <a href="#">AMBARI-18093</a> | 2.2.2                | 2.4.0       | <p>On large clusters, if the TTL of high precision tables is more than 3 days, it leads to too much data and too many regions in AMS HBase. It is better to have a smaller ttl for the higher precision data. The 5 minute aggregate data will still be available for 7 days to work with.</p>  | <ol style="list-style-type: none"> <li>Following config changes in "ams-site" from the UI<br/> <code>timeline.metrics.host.aggregator.ttl : 1 day</code><br/> <code>timeline.metrics.cluster.aggregator.second.ttl : 3 days</code></li> <li>Restart AMS collector</li> </ol>   |

|    |                              |                 |       |   |  |
|----|------------------------------|-----------------|-------|---|--|
| 11 | <a href="#">AMBARI-1779</a>  | 2.2.1, 2.2.2    | 2.4.0 | <p>The HBase normalizer which automatically splits / merges regions based on region size was leveraged in AMS (2.2.1). However, due to over aggressive region splitting by normalizer sometimes, in large clusters it could lead to an explosion of regions. This will eventually lead to AMS crashing every time it starts up.</p> <p>As of 2.2.2, the AMS HBase normalizer cannot be disabled through AMS configs.</p>          | <p>Instructions for disabling normalizer on AMS HBase tables.</p> <ol style="list-style-type: none"> <li>1. su ams (kinit if needed)</li> <li>2. Connect to HBase shell on Metrics collector host.<br/><code>/usr/lib/ams-hbase/bin/hbase -config /etc/ams-hbase/conf shell</code></li> <li>3. Execute the following statements one by one.</li> </ol> <pre> alter 'METRIC_RECORD', {NORMALIZATION_ENABLED =&gt; 'false'} alter 'METRIC_AGGREGATE', {NORMALIZATION_ENABLED =&gt; 'false'} alter 'METRIC_RECORD_MINUTE', {NORMALIZATION_ENABLED =&gt; 'false'} alter 'METRIC_AGGREGATE_MINUTE', {NORMALIZATION_ENABLED =&gt; 'false'} alter 'METRIC_RECORD_HOURLY', {NORMALIZATION_ENABLED =&gt; 'false'} alter 'METRIC_AGGREGATE_HOURLY', {NORMALIZATION_ENABLED =&gt; 'false'} alter 'METRIC_RECORD_DAILY', {NORMALIZATION_ENABLED =&gt; 'false'} alter 'METRIC_AGGREGATE_DAILY', {NORMALIZATION_ENABLED =&gt; 'false'} </pre> <p>4. Verify the configuration change took effect. Open the HBase master UI on a browser (<code>http://&lt;collector_host&gt;:61310</code>), and search for the String "NORMALIZATION". It should return no matches.</p> |
| 10 | <a href="#">AMBARI-15492</a> | 2.2.1           | 2.2.2 | <p>Ambari metrics collector shuts down and restarts randomly. The following error is seen in the collector log.</p> <p><code>ERROR org.apache.hadoop.yarn.server.applicationhistoryservice. ApplicationHistoryServer: RECEIVED SIGNAL 15: SIGTERM</code></p>  | <ul style="list-style-type: none"> <li>Comment out these 2 properties in <code>/etc/ambari-server/conf/ambari.properties</code> <ul style="list-style-type: none"> <li><code>#recovery.enabled_components=METRICS_COLLECTOR</code></li> <li><code>#recovery.type=AUTO_START</code></li> </ul> </li> <li>Restart Ambari Server.</li> </ul>  |
| 9  | <a href="#">AMBARI-13758</a> | 2.1.2 and lower | 2.2.0 | When Ambari Metrics collector host is moved from one host to another, host metrics are not seen   | <ul style="list-style-type: none"> <li>Restart Ambari Server.</li> <li>Restart Ambari Metrics Monitor</li> </ul>   |
| 8  | <a href="#">AMBARI-14257</a> | 2.1.2           | 2.2.0 | Storm metrics are not seen after upgrading to Ambari 2.1.2  | <p>On every host with a storm component [nimbus / supervisor / client], carry out the following steps.</p> <ol style="list-style-type: none"> <li>1. Verify broken link<br/><code>ls -al /usr/hdp/current/storm-&lt;component&gt;/lib/ ambari-metrics-storm-sink.jar</code><br/>Remove symlink</li> <li>2. <code>rm -f /usr/hdp/current/storm-&lt;component&gt;/lib/ ambari-metrics-storm-sink.jar</code><br/>Reattach symlink to new JAR</li> <li>3. <code>ln -s /usr/lib/storm/lib/ambari-metrics-storm- sink*.jar /usr/hdp/current/storm-&lt;component&gt;/lib/ambari- metrics-storm-sink.jar</code></li> </ol>   |
| 7  | <a href="#">AMBARI-13798</a> | 2.1.2           | 2.2.0 | <p>Ambari Metrics service graphs might not show data for certain metrics and the the following error might be seen in the metrics collector log (Ambari 2.1.2, 2.1.2.1).</p> <p><code>"The time range query for precision table exceeds row count limit, please query aggregate table instead"</code></p>   | <ul style="list-style-type: none"> <li>Change the ams-site configuration, Set <code>timeline.metrics.service.default.result.limit = 15840</code></li> <li>Restart the Collector</li> </ul>   |
| 6  | <a href="#">AMBARI-13711</a> | 2.1.2           |       | <p>Ambari Metrics Server wont start successfully with Kerberos in Distributed Mode (<a href="#">AMBI-13711</a>)</p> <p>The problem is, we cannot have separate principals for HBase Master and RS. The Zookeeper ACLs will not allow znode created using one principal to be read by the other unless proper ACL are set.</p> <p>Since master created the znode with a different principal than RS in 2.1.2 this will happen.</p> | <p>Change the AMS configuration to use the Master keytab and principal for RS Set,</p> <ul style="list-style-type: none"> <li><code>ams-hbase-site :: hbase.regionserver.keytab.file = /etc/security/keytabs/ams-hbase.master.keytab</code></li> <li><code>ams-hbase-site :: hbase.regionserver.kerberos.principal = amshbasemaster/_HOST@REALM</code></li> </ul> <p>Restart the Collector</p>   |
| 5  |                              |                 |       | Metrics data for last one month is missing many data points that should exist.  | <ul style="list-style-type: none"> <li>- check <code>/var/log/ambari-metrics-collector/ambari-metrics-collector.log</code> for metrics data aggregation errors like <code>OutOfOrderScannerNextException</code> or <code>SpoolTooBigToDiskException</code></li> <li>- Set bigger value for <code>hbase_regionserver_heapsize</code> property in Advanced ams-hbase-env using Ambari Web UI</li> <li>- Restart Metrics Collector.</li> </ul>  |

|   |                               |                     |       |   |  |
|---|-------------------------------|---------------------|-------|---|--|
| 4 | AMBARI-11501,<br>AMBARI-12347 | 2.0.x and 2.1.0     |       | AMS HBase does not start after Kerberization in distributed mode, 2.0.x and 2.1.0<br>(Note: Also look at issue: 1)                                  | <p>Issue 1: (<a href="#">AMBARI-11501</a>)<br/><i>Steps to workaround this issue:</i></p> <p>On the ambari server host</p> <ul style="list-style-type: none"> <li>• cd /var/lib/ambari-agent/cache/common-services/AMBARI_METRICS/0.1.0/package/scripts/params.py</li> <li>• Edit <code>hbase_staging_dir</code> and point it to the desired location on HDFS, e.g.: "/ams-hbase/staging"</li> <li>• Restart Ambari server.</li> <li>• Restart Metrics Collector.</li> </ul> <p>Issue 2: (<a href="#">AMBARI-12347</a>)<br/><i>Steps to workaround this issue:</i></p> <ul style="list-style-type: none"> <li>• Edit the following properties in <a href="#">Advanced.ams-hbase-security-site</a> using Ambari Web UI:</li> </ul> <pre>ams.zookeeper.principal = <code>zookeeper/_HOST@EXAMPLE.COM</code><br/>(Substitute appropriate REALM)</pre> <pre>ams.zookeeper.keytab = /etc/security/keytabs/zk.service.keytab</pre> <p><b>Note:</b> This is assuming you have a Zookeeper keytab on the host with Metrics Collector.</p> <p>If not you should create one with appropriate permissions.</p> <p>If a keytab already exists, make sure to <code>chmod 440 /etc/security/keytabs/zk.service.keytab</code></p> <p>Example: <code>l# klist -kt /etc/security/keytabs/zk.service.keytab</code></p> <pre>Keytab name: FILE:/etc/security/keytabs/zk.service.keytab KVNO Timestamp Principal -----<br/>-----<br/>1 07/08/15 22:29:07 zookeeper/ambari-sid-3.c.pramod-thangali.internal@EXAMPLE.COM</pre> <ul style="list-style-type: none"> <li>• Restart Metrics Collector</li> </ul> |
| 3 |                               | 2.0.0, 2.1.0, 2.1.1 | 2.1.2 | Alter TTL is not supported by the version of Phoenix used with Ambari 2.0.0, 2.1.0, 2.1.1. This property can be modified using HBase Shell command. | <pre>-]\$ su - ams -]\$ export JAVA_HOME=/usr/jdk64/jdk1.8.0_40/ -]\$ /usr/lib/ams-hbase/bin/hbase --config /etc/ams-hbase/conf shell  hbase(main):007:0&gt; describe 'METRIC_RECORD' ----- HBase output describing table information -----  hbase(main):009:0&gt; alter 'METRIC_RECORD', { NAME =&gt; '0', TTL =&gt; 172800}  hbase(main):007:0&gt; describe 'METRIC_RECORD' ----- HBase output describing table information with new TTL -----</pre>   |
| 2 |                               | 2.0.x               | 2.1.0 | Ambari Metrics service does not work after enabling security with AMS in distributed mode in Ambari 2.0.x   | <ol style="list-style-type: none"> <li>1. Copy core-site.xml and hdfs-site.xml to metric collector config directories after HA is enabled.</li> </ol> <pre>cp /etc/hadoop/conf/core-site.xml /etc/ambari-metrics-collector/conf/ cp /etc/hadoop/conf/hdfs-site.xml /etc/ambari-metrics-collector/conf/ cp /etc/hadoop/conf/core-site.xml /etc/ams-hbase/conf/ cp /etc/hadoop/conf/hdfs-site.xml /etc/ams-hbase/conf/</pre> <ol style="list-style-type: none"> <li>2. Restart the Metrics Collector Component.</li> </ol>   |
| 1 | AMBARI-10707                  | 2.0.x               | 2.1.0 | Ambari Metrics service does not work with NN HA in distributed mode in Ambari 2.0.x   |  |