


HIVE-9453.2.patch #2

Edit

 Open

 thejasmn wants to merge 1 commit into trunk from HIVE-9453

 Conversation 8

 Commits 1

 Files changed 46

+8,340 -42 




thejasmn commented 2 days ago

Owner 

No description provided.


 HIVE-9453.2.patch 28a6b82

 thejasmn commented on the diff a day ago

.../apache/hadoop/hive/metastore/WritableMicroBench.java [View full changes](#)




((17 lines not shown))
17 +import java.io.ObjectInputStream;
18 +import java.io.ObjectOutputStream;
19 +import java.nio.charset.StandardCharsets;
20 +import java.security.MessageDigest;
21 +import java.util.ArrayList;
22 +import java.util.Arrays;
23 +import java.util.HashMap;
24 +import java.util.HashSet;
25 +import java.util.List;
26 +import java.util.Map;
27 +import java.util.Set;
28 +
29 +/**
30 + * Created by gates on 12/17/14.
31 + */
32 +public class WritableMicroBench {

 thejasmn added a note a day ago


Owner  

Is this file to be committed ?


Add a line note

Labels 

None yet


Milestone 

No milestone

Assignee 

No one—assign yourself

Notifications


 Unsubscribe

You're receiving notifications because you authored the thread.


1 participant



 Lock pull request

 thejasmn commented on the diff a day ago

...che/hadoop/hive/metastore/hbase/DatabaseWritable.java [View full changes](#)



((14 lines not shown))
14 + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY
15 + * KIND, either express or implied. See the License for the
16 + * specific language governing permissions and limitations
17 + * under the License.
18 + */
19 +package org.apache.hadoop.hive.metastore.hbase;
20 +
21 +import org.apache.hadoop.hive.metastore.api.Database;
22 +import org.apache.hadoop.io.Writable;
23 +
24 +import java.io.DataInput;

```
25 +import java.io.DataOutput;  
26 +import java.io.IOException;  
27 +  
28 +/**  
29 + * Wrapper for {@link org.apache.hadoop.hive.metastore.api.Table} that makes it w
```



thejasmn added a note a day ago

Owner



Wrapper for api.Database (not api.Table)

Add a line note

<> thejasmn commented on the diff 23 hours ago

...g/apache/hadoop/hive/metastore/hbase/ObjectCache.java

[View full changes](#)

((12 lines not shown))

```
12 + * Unless required by applicable law or agreed to in writing,  
13 + * software distributed under the License is distributed on an  
14 + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY  
15 + * KIND, either express or implied. See the License for the  
16 + * specific language governing permissions and limitations  
17 + * under the License.  
18 + */  
19 +package org.apache.hadoop.hive.metastore.hbase;  
20 +  
21 +import java.util.HashMap;  
22 +import java.util.Map;  
23 +  
24 +/**  
25 + * A generic class for caching objects obtained from HBase. Currently a set of  
26 + * convenience methods around a {@link java.util.HashMap} with a max size but built  
27 + * as a separate class in case we want to switch out the implementation to something
```



thejasmn added a note 23 hours ago

Owner



a comment saying that when max limit is hit, addition to the cache is ignored would be useful.



thejasmn added a note 21 hours ago

Owner



Can you also add comment saying that this cache is not synchronized because it is going to be used from only one thread (it is indirectly threadlocal)

Add a line note

<> thejasmn commented on the diff 21 hours ago

...cache/hadoop/hive/metastore/hbase/PartitionCache.java

[View full changes](#)

((17 lines not shown))

```
17 + * under the License.  
18 + */  
19 +package org.apache.hadoop.hive.metastore.hbase;  
20 +  
21 +import org.apache.hadoop.hive.common.ObjectPair;  
22 +import org.apache.hadoop.hive.metastore.api.Partition;  
23 +  
24 +import java.util.Collection;  
25 +import java.util.HashMap;  
26 +import java.util.List;  
27 +import java.util.Map;  
28 +  
29 +/**  
30 + * A cache for partition objects. This is separate from ObjectCache because we ne
```

```
31 + * differently (always by table) and because we need to be able to track whether v
32 + * all of the partitions for a table or not.
```



thejasmn added a note 21 hours ago

Owner



a comment mentioning that this is to be used only from a single thread would be useful.

[Add a line note](#)

<> thejasmn commented on the diff 3 hours ago

...rg/apache/hadoop/hive/metastore/hbase/StatsCache.java

[View full changes](#)

((176 lines not shown))

```
176 + // if the indicated key is already gone.
177 + cache.remove(key);
178 + } finally {
179 + lock.writeLock().unlock();
180 + }
181 + }
182 +
183 + private void clean() {
184 + // This spawns a separate thread to walk through the cache and clean.
185 + synchronized (this) {
186 + if (cleaning) return;
187 + cleaning = true;
188 + cleans.incr();
189 + }
190 +
191 + Thread cleaner = new Thread() {
```



thejasmn added a note 3 hours ago

Owner



Indentation is off

[Add a line note](#)

<> thejasmn commented on the diff 43 minutes ago

...rg/apache/hadoop/hive/metastore/hbase/StatsCache.java

[View full changes](#)

((93 lines not shown))

```
93 + counters.add(cleans);
94 + }
95 +
96 + /**
97 + * Add an object to the cache.
98 + * @param dbName name of database table is in
99 + * @param tableName name of table
100 + * @param partName name of partition, can be null if these are table level stats
101 + * @param colName name of the column these statistics are for
102 + * @param stats stats
103 + * @param lastAnalyzed last time these stats were analyzed
104 + */
105 + void put(String dbName, String tableName, String partName, String colName,
106 + ColumnStatisticsObj stats, long lastAnalyzed) {
107 + if (cache.size() >= maxSize) clean();
108 + Key key = new Key(dbName, tableName, partName, colName);
```




thejasmn added a note 43 minutes ago

Owner



should we avoid adding this key to cache if it is currently at the limit ? It is not certain that clean() call is going to free up any entries (all of them could return false for tooLate)

Add a line note

<>  thejasmn commented on the diff 38 minutes ago

...rg/apache/hadoop/hive/metastore/hbase/StatsCache.java

[View full changes](#)



((201 lines not shown))

```
201 +         entries = new ArrayList<Map.Entry<Key, StatsInfo>>(cache.entrySet());
202 +     } finally {
203 +         lock.readLock().unlock();
204 +     }
205 +     for (Map.Entry<Key, StatsInfo> entry : entries) {
206 +         if (tooLate(entry.getValue())) {
207 +             remove(entry.getKey());
208 +         }
209 +         // We want to make sure this runs at a low priority in the background
210 +         Thread.yield();
211 +     }
212 + } catch (Throwable t) {
213 +     // Don't let anything past this thread that could end up killing the r
214 +     LOG.error("Caught exception in stats cleaner thread", t);
215 + } finally {
216 +     cleaning = false;
```



thejasmn added a note 38 minutes ago

Owner



should we check if were able to remove entries before considering the cleaning to be done? In cases where clean cycle does not remove entries, we could dynamically adjust the value of timeToLive (for example based on min lastTouch value) and run another clean cycle.

Add a line note

Add more commits by pushing to the **HIVE-9453** branch on **thejasmn/hive**.



This pull request can be automatically merged.

You can also merge branches on the [command line](#).



Merge pull request



Write

Preview

Markdown supported [Edit in fullscreen](#)

Leave a comment

Attach images by dragging & dropping, [selecting them](#), or pasting from the clipboard.

Close pull request

Comment

ProTip! Add comments to specific lines under [Files changed](#).

