

Kune - Defect #525

SiteMap generation is not working properly in large sites

03/22/2013 10:36 AM - Vicente J. Ruiz Jurado

Status:	New	Start date:	03/22/2013
Priority:	High	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Tags:	
Resolution:			
Description			
The current exception			
<pre>2013-03-22 00:20:00 GenerateSitemapJob [INFO] Daily sitemap generation cron job start 2013-03-22 00:20:00 LoggerMethodInterceptor [DEBUG] SiteMapGenerator.generate() 2013-03-22 00:20:00 NewPooledConnection [WARN] [c3p0] A PooledConnection that has already signalled a Connection error is still in use! 2013-03-22 00:20:00 NewPooledConnection [WARN] [c3p0] Another error has occurred [com.mysql.jdbc.exceptions.jdbc4.MySQLNonTransientConnectionException: No operations allowed after connection closed.] which will not be reported to listeners! com.mysql.jdbc.exceptions.jdbc4.MySQLNonTransientConnectionException: No operations allowed after connection closed. at sun.reflect.NativeConstructorAccessorImpl.newInstance0(Native Method) at sun.reflect.NativeConstructorAccessorImpl.newInstance(NativeConstructorAccessorImpl.java:57) at sun.reflect.DelegatingConstructorAccessorImpl.newInstance(DelegatingConstructorAccessorImpl.java:45) at java.lang.reflect.Constructor.newInstance(Constructor.java:532) at com.mysql.jdbc.Util.handleNewInstance(Util.java:409) at com.mysql.jdbc.Util.getInstance(Util.java:384) at com.mysql.jdbc.SQLError.createSQLException(SQLError.java:1015) at com.mysql.jdbc.SQLError.createSQLException(SQLError.java:989) at com.mysql.jdbc.SQLError.createSQLException(SQLError.java:984) at com.mysql.jdbc.SQLError.createSQLException(SQLError.java:929) at com.mysql.jdbc.ConnectionImpl.throwConnectionClosedException(ConnectionImpl.java:1193) at com.mysql.jdbc.ConnectionImpl.checkClosed(ConnectionImpl.java:1180) at com.mysql.jdbc.ConnectionImpl.prepareStatement(ConnectionImpl.java:4137) at com.mysql.jdbc.ConnectionImpl.prepareStatement(ConnectionImpl.java:4103) at com.mchange.v2.c3p0.impl.NewProxyConnection.prepareStatement(NewProxyConnection.java:213) at sun.reflect.GeneratedMethodAccessor12.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:616) at org.hibernate.engine.jdbc.internal.proxy.ConnectionProxyHandler.continueInvocation(ConnectionProxyHandler.java:138) at org.hibernate.engine.jdbc.internal.proxy.AbstractProxyHandler.invoke(AbstractProxyHandler.java:81) at \$Proxy100.prepareStatement(Unknown Source) at org.hibernate.engine.jdbc.internal.StatementPreparerImpl\$5.doPrepare(StatementPreparerImpl.java:147) at org.hibernate.engine.jdbc.internal.StatementPreparerImpl\$StatementPreparationTemplate.prepareStatement(StatementPreparerImpl.java:166) at org.hibernate.engine.jdbc.internal.StatementPreparerImpl.prepareQueryStatement(StatementPreparerImpl.java:145) at org.hibernate.loader.Loader.prepareQueryStatement(Loader.java:1720) at org.hibernate.loader.Loader.doQuery(Loader.java:828) at org.hibernate.loader.Loader.doQueryAndInitializeNonLazyCollections(Loader.java:289) at org.hibernate.loader.Loader.doList(Loader.java:2447) at org.hibernate.loader.Loader.doList(Loader.java:2433)</pre>			

```
at org.hibernate.loader.Loader.listIgnoreQueryCache(Loader.java:2263)
at org.hibernate.loader.Loader.list(Loader.java:2258)
at org.hibernate.loader.hql.QueryLoader.list(QueryLoader.java:470)
at org.hibernate.hql.internal.ast.QueryTranslatorImpl.list(QueryTranslatorImpl.java:355)
at org.hibernate.engine.query.spi.HQLQueryPlan.performList(HQLQueryPlan.java:195)
at org.hibernate.internal.SessionImpl.list(SessionImpl.java:1215)
at org.hibernate.internal.QueryImpl.list(QueryImpl.java:101)
at org.hibernate.ejb.QueryImpl.getSingleResult(QueryImpl.java:284)
at com.google.inject.persist.jpa.JpaFinderProxy.invoke(JpaFinderProxy.java:72)
at com.google.inject.persist.jpa.JpaPersistModule$1.invoke(JpaPersistModule.java:129)
at $Proxy23.countExceptType(Unknown Source)
at cc.kune.core.server.searchable.SiteMapGenerator.countNotClosedGroups(SiteMapGenerator.java:160)
at cc.kune.core.server.searchable.SiteMapGenerator.generate(SiteMapGenerator.java:115)
at cc.kune.core.server.searchable.SiteMapGenerator$$EnhancerByGuice$$26e8b0a.CGLIB$generate$0(<generated>)
at
cc.kune.core.server.searchable.SiteMapGenerator$$EnhancerByGuice$$26e8b0a$$FastClassByGuice$$a56540aa.invoke(<generated>
ted)
```

Related issues:

Related to Enhancement # 520: Launch sitemap generation from MBean	Closed	03/12/2013
Related to Enhancement # 286: Make kune googleable/searchable using hash bang...	In Progress	03/04/2013

Associated revisions

Revision ed0ca096 - 03/01/2013 12:56 PM - Vicente J. Ruiz Jurado

Added some more indexes to db (thanks to Pablo Ojanguren @pablojan). Fix #510

Revision 6fe46139 - 03/13/2013 04:51 PM - Vicente J. Ruiz Jurado

Sitemap transaction added of group counts. Continuation of #286

Revision af63bbbe - 03/22/2013 11:02 AM - Vicente J. Ruiz Jurado

Some test for #525

Revision 28356438 - 03/27/2013 09:51 PM - Vicente J. Ruiz Jurado

Some work with #520 #525 and following #70

Revision a3a1664e - 03/27/2013 09:55 PM - Vicente J. Ruiz Jurado

Some work with #525

History

#1 - 03/26/2013 11:21 AM - Pablo Ojanguren

It looks like some query is taking too much time and then, pooled connection expires.

There are 3 SQL queries in the algorithm. I am looking into the final SQL executed for every one.

- Query for Method SiteMapGenerator.getGroups()

```
select group0_.id as id7_, group0_.admissionType as admission7_, group0_.backgroundImage as background7_, group0_.backgroundMime as background4_7_, group0_.createdOn as createdOn7_, group0_.defaultContent_id as default14_7_, group0_.defaultLicense_id as default15_7_, group0_.groupType as groupType7_, group0_.logo as logo7_, group0_.logoLastModifiedTime as logoLast8_7_, group0_.mimesubtype as mimesubt9_7_, group0_.mimetype as mimetype7_, group0_.longName as longName7_, group0_.shortName as shortName7_, group0_.socialNetwork_id as socialN16_7_, group0_.workspaceTheme as workspace13_7_ from groups group0_ where group0_.groupType<>1 order by group0_.createdOn DESC limit ?;
```

Executing the query w/o limit from a mySql client, it took 0.265 sec / 77.555 sec (execute / fetch)

- Query 2

```
select container0_.id as id1_, container0_.accessLists_id as accessLi7_1_, container0_.createdOn as createdOn1_, container0_.deletedOn as deletedOn1_, container0_.language_id as language8_1_, container0_.name as name1_, container0_.owner_id as owner9_1_, container0_.parent_id as parent10_1_, container0_.toolName as toolName1_, container0_.typeid as typeid1_ from containers container0_ where $GROUP_ID in (select group6_.id from groups group1_ inner join social_networks socialnetw2_ on group1_.socialNetwork_id=socialnetw2_.id inner join access_lists accesslist3_ on socialnetw2_.accessLists_id=accesslist3_.id inner join group_list grouplist4_ on accesslist3_.editors_id=grouplist4_.id, group_list_groups list5_, groups group6_ where container0_.owner_id=group1_.id and grouplist4_.id=list5_.group_list_id and list5_.list_id=group6_.id) or $GROUP_ID in (select group12_.id from groups group7_ inner join social_networks socialnetw8_ on group7_.socialNetwork_id=socialnetw8_.id inner join access_lists accesslist9_ on socialnetw8_.accessLists_id=accesslist9_.id inner join group_list grouplist10_ on accesslist9_.admins_id=grouplist10_.id, group_list_groups list11_, groups group12_ where container0_.owner_id=group7_.id and grouplist10_.id=list11_.group_list_id and list11_.list_id=group12_.id);
```

- Query 3

```
select content0_.id as id3_, content0_.accessLists_id as accessL13_3_, content0_.container_id as container14_3_, content0_.createdOn as createdOn3_, content0_.deletedOn as deletedOn3_, content0_.filename as filename3_, content0_.language_id as language15_3_, content0_.lastRevision_id as lastRev16_3_, content0_.license_id as license17_3_, content0_.mimesubtype as mimesubt5_3_, content0_.mimetype as mimetype3_, content0_.modifiedOn as modifiedOn3_, content0_.publishedOn as published8_3_, content0_.status as status3_, content0_.typeid as typeid3_, content0_.version as version3_, content0_.waveId as waveId3_ from contents content0_ where $GROUP_ID in (select group7_.id from containers container1_ inner join groups group2_ on container1_.owner_id=group2_.id inner join social_networks socialnetw3_ on group2_.socialNetwork_id=socialnetw3_.id inner join access_lists accesslist4_ on socialnetw3_.accessLists_id=accesslist4_.id inner join group_list grouplist5_ on accesslist4_.editors_id=grouplist5_.id, group_list_groups list6_, groups group7_ where content0_.container_id=container1_.id and grouplist5_.id=list6_.group_list_id and list6_.list_id=group7_.id) or $GROUP_ID in (select group14_.id from containers container8_ inner join groups group9_ on container8_.owner_id=group9_.id inner join social_networks socialnetw10_ on group9_.socialNetwork_id=socialnetw10_.id inner join access_lists accesslist11_ on socialnetw10_.accessLists_id=accesslist11_.id inner join group_list grouplist12_ on accesslist11_.admins_id=grouplist12_.id, group_list_groups list13_, groups group14_ where content0_.container_id=container8_.id and grouplist12_.id=list13_.group_list_id and list13_.list_id=group14_.id) order by content0_.modifiedOn DESC;
```

Initially it doesn't look a long-last query problem. I've executed separately those queries for a sample of data input (groups) and they didn't take more than a couple of seconds.

We need to analyze Java code.

Something weird is that the exception is in the first (SiteMapGenerator.java:160).

#3 - 04/01/2013 09:56 AM - Pablo Ojanguren

So problem comes from a SELECT COUNT... (SiteMapGenerator.countNotClosedGroups(SiteMapGenerator.java:160))

We could start thinking on hibernate-pool configuration problem.

#4 - 04/02/2013 09:43 AM - Pablo Ojanguren

After today's comment, we could think that there is any problem with these queries. Let's suppose that system's performance cause the DB time out.

We should make more tests, increasing the time out value for the c3p0 pool.