

D4Science Infrastructure - Task #4617

Smart-executor- node24.d4science.org: Too Many Open Files

Jul 04, 2016 10:20 AM - Roberto Cirillo

Status:	Closed	Start date:	Jul 04, 2016
Priority:	Normal	Due date:	
Assignee:	Costantino Perciante	% Done:	100%
Category:	Data Management	Estimated time:	0.00 hour
Target version:	UnSprintable		
Infrastructure:	Production		

Description

On node24.d4science.org is running the social-data-indexer plugin of SmartExecutor service.

We have the following exceptions:

catalina.out:

```
java.lang.NullPointerException
com.netflix.astyanax.connectionpool.exceptions.PoolTimeoutException: PoolTimeoutException: [host=cassandra2-p-d4s.d4science.org(146.48.123.140):9160, latency=6000(6000), attempts=3]Timed out waiting for connection
    at com.netflix.astyanax.connectionpool.impl.SimpleHostConnectionPool.waitForConnection(SimpleHostConnectionPool.java:218)
    at com.netflix.astyanax.connectionpool.impl.SimpleHostConnectionPool.borrowConnection(SimpleHostConnectionPool.java:185)
    at com.netflix.astyanax.connectionpool.impl.RoundRobinExecuteWithFailover.borrowConnection(RoundRobinExecuteWithFailover.java:66)
    at com.netflix.astyanax.connectionpool.impl.AbstractExecuteWithFailoverImpl.tryOperation(AbstractExecuteWithFailoverImpl.java:67)
    at com.netflix.astyanax.connectionpool.impl.AbstractHostPartitionConnectionPool.executeWithFailover(AbstractHostPartitionConnectionPool.java:253)
    at com.netflix.astyanax.thrift.ThriftClusterImpl.describeKeyspaces(ThriftClusterImpl.java:155)
    at com.netflix.astyanax.thrift.ThriftClusterImpl.describeKeyspace(ThriftClusterImpl.java:174)
    at org.gcube.portal.databook.server.CassandraClusterConnection.SetupKeySpaces(CassandraClusterConnection.java:157)
    at org.gcube.portal.databook.server.CassandraClusterConnection.<init>(CassandraClusterConnection.java:101)
    at org.gcube.portal.databook.server.DBCassandraAstyanaxImpl.<init>(DBCassandraAstyanaxImpl.java:201)
    at org.gcube.socialnetworking.socialdataindexer.SocialDataIndexerPlugin.launch(SocialDataIndexerPlugin.java:95)
    at org.gcube.vremanagement.executor.pluginmanager.RunnablePlugin.run(RunnablePlugin.java:67)
    at org.gcube.vremanagement.executor.scheduler.SmartExecutorTask.execute(SmartExecutorTask.java:214)
    at org.quartz.core.JobRunShell.run(JobRunShell.java:202)
    at org.quartz.simpl.SimpleThreadPool$WorkerThread.run(SimpleThreadPool.java:573)
java.lang.NullPointerException
```

ghn.log

```
00:00:11.560 [pool-2-thread-1] WARN ProfileBuilder: unable to detect the uptime of this machine
java.io.IOException: Cannot run program "uptime": error=24, Too many open files
    at java.lang.ProcessBuilder.start(ProcessBuilder.java:1047) ~[na:1.7.0_80]
    at java.lang.Runtime.exec(Runtime.java:617) ~[na:1.7.0_80]
    at java.lang.Runtime.exec(Runtime.java:450) ~[na:1.7.0_80]
    at java.lang.Runtime.exec(Runtime.java:347) ~[na:1.7.0_80]
    at org.gcube.smartgears.handlers.container.lifecycle.ProfileBuilder.uptime(ProfileBuilder.java:297) [common-smartgears-1.2.7-3.11.0-128702.jar:na]
    at org.gcube.smartgears.handlers.container.lifecycle.ProfileBuilder.update(ProfileBuilder.java:228) [common-smartgears-1.2.7-3.11.0-128702.jar:na]
```

```
at org.gcube.smartgears.handlers.container.lifecycle.ProfileManager$2$1.run(ProfileManager
.java:266) [common-smartgears-1.2.7-3.11.0-128702.jar:na]
at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:471) [na:1.7.0_80]
at java.util.concurrent.FutureTask.runAndReset(FutureTask.java:304) [na:1.7.0_80]
at java.util.concurrent.ScheduledThreadPoolExecutor$ScheduledFutureTask.access$301(Schedul
edThreadPoolExecutor.java:178) [na:1.7.0_80]
at java.util.concurrent.ScheduledThreadPoolExecutor$ScheduledFutureTask.run(ScheduledThrea
dPoolExecutor.java:293) [na:1.7.0_80]
at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145) [na:1.7
.0_80]
at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615) [na:1.7
.0_80]
at java.lang.Thread.run(Thread.java:745) [na:1.7.0_80]
Caused by: java.io.IOException: error=24, Too many open files
at java.lang.UNIXProcess.forkAndExec(Native Method) ~[na:1.7.0_80]
at java.lang.UNIXProcess.<init>(UNIXProcess.java:187) ~[na:1.7.0_80]
at java.lang.ProcessImpl.start(ProcessImpl.java:130) ~[na:1.7.0_80]
at java.lang.ProcessBuilder.start(ProcessBuilder.java:1028) ~[na:1.7.0_80]
... 13 common frames omitted
```

The plugin was not working and I've restarted the container. It's need further analysis.

Related issues:

Related to D4Science Infrastructure - Task #4647: node24.d4science.org : Incr...

Closed

Jul 06, 2016

History

#1 - Jul 06, 2016 11:23 AM - Roberto Cirillo

- Related to Task #4647: node24.d4science.org : Increase the maximum number of file descriptors added

#2 - Jul 11, 2016 10:20 AM - Costantino Perciante

- Status changed from New to In Progress

- % Done changed from 0 to 90

The problem arises from the fact that I used two libraries in the plugin, namely the social networking library and the elastic search client library that do not automatically close the connection pool they open. I've just modified the first library and I'm going to test the it during this week to check that the close works as it should. As far as the second library, I've already tested its close mechanism and it works. I will update this ticket as soon as I'm sure everything works fine. After that, we can switch back and reduce the number of file descriptors that can be opened.

#3 - Sep 27, 2016 10:34 AM - Costantino Perciante

- Status changed from In Progress to Feedback

- % Done changed from 90 to 100

The new social-data-indexer-plugin which is going to be released in gcube 4.1 works as expected. Connections pool are closed and the number of connections used during the run are at most 20, so I'm going to put this ticket to feedback. You can decrease back the number of allowed sockets after 4.1 goes to production.

#4 - Nov 17, 2016 06:04 PM - Roberto Cirillo

- Status changed from Feedback to Closed