

D4Science Infrastructure - Task #9551

Some jobs wrote big chunks of data into the persistence jackrabbit table

Aug 30, 2017 11:32 PM - Andrea Dell'Amico

Status:	In Progress	Start date:	Aug 30, 2017
Priority:	Normal	Due date:	
Assignee:	Lucio Lelii	% Done:	10%
Category:	System Application	Estimated time:	0.00 hour
Target version:	Production Jackrabbit Migration from Derby to PostgreSQL		
Infrastructure:	Production		
Description			
From Aug 19th and 21th something (someone) wrote repeatedly into the workspace big chunks of data into the persistence manager data field.			
We need to investigate and find the rogue process to avoid a similar event in the future.			

History

#2 - Sep 01, 2017 11:02 AM - Valentina Marioli

- Status changed from New to In Progress

#4 - Sep 01, 2017 11:47 AM - Gianpaolo Coro

When a computation fails, a gCube Item is created having a stacktrace attached. In the case of BiOnym, this can be long since it contains the logs by F. Fiorellato's (ex-FAO) libraries for taxonomic names matching. If the WS or other system was blocked for some reason, the ~3400 calls per day could have reported a large amount of stacktrace logs. Thus, the large number of logs and occupied space could be an indicator that either the WS system or other satellite services could be not working. For example, if the URI-Resolver was unavailable for some reason, BiOnym would have failed all its computations (because it uses files accessed on the WS through the URI-Resolver).

Some questions need clarification:

- 1 - If either the WS or satellite services were not working, did we have nagios alerts?
- 2 - Were there issues on the URI-Resolver in those days?
- 3 - Are the logs in the BiOnym stacktrace so huge to justify that occupation?

#5 - Sep 01, 2017 12:09 PM - Andrea Dell'Amico

I also uploaded a blob of something that seems an execution log, with failures: <https://goo.gl/VNmpeJ>

#6 - Sep 01, 2017 12:15 PM - Andrea Dell'Amico

Gianpaolo Coro wrote:

When a computation fails, a gCube Item is created having a stacktrace attached. In the case of BiOnym, this can be long since it contains the logs by F. Fiorellato's (ex-FAO) libraries for taxonomic names matching. If the WS or other system was blocked for some reason, the ~3400 calls per day could have reported a large amount of stacktrace logs. Thus, the large number of logs and occupied space could be an indicator that either the WS system or other satellite services could be not working. For example, if the URI-Resolver was unavailable for some reason, BiOnym would have failed all its computations (because it uses files accessed on the WS through the URI-Resolver).

Some questions need clarification:

- 1 - If either the WS or satellite services were not working, did we have nagios alerts?

Yes. The infra portals, workspace, mongo and the URI resolver are monitored. The URI resolver is also under HA since some months.

- 2 - Were there issues on the URI-Resolver in those days?

No

- 3 - Are the logs in the BiOnym stacktrace so huge to justify that occupation?

Those blobs do not contain stacktraces only. The second one I posted it's a stack trace, the first one if full of the same strings repeated thousands of times.

#7 - Sep 01, 2017 01:02 PM - Costantino Perciante

If we really want to understand how the content is stored in there, then we will need to inspect the jackrabbit code. I think it would be a quite expensive operation.

Andrea, I think we should restore (somewhere else of course) a dump of that database (we can discuss about the dates.. something before 22-23/08 can be useful). This will allow us to compare some rows of the pm_default_bundle table in terms of size and content (and presence as well)

Also, we figured out that postgres (luckily) compresses data, so octet_length when a bytea field is involved just returns the size of the compressed data (the biggest rows needs 23MB in postgres, but after decompression it takes what Andrea says.. i.e. more than 400 MB)

#8 - Sep 01, 2017 01:10 PM - Andrea Dell'Amico

@valentina.marioli@isti.cnr.it just asked to restore those backups. You can find them on workspace-repository-prod1.d4science.org under /data. There's a backup from August 19th and one from Aug 16th.

#9 - Sep 01, 2017 02:46 PM - Gianpaolo Coro

One of the big blobs contains a large quantity of objects generated by the old-good Statistical Manager between 2014 and Nov. 2016. These files were written in the ".Application" hidden folder, which is empty for the statistical.manager user. Since this logic was contained in the StatMan service (not in the algorithms execution engine), it is not contained in any way in the DataMiner. Those are old objects that have come from somewhere. Valentina is now checking if they were present also in the Derby DB. Further, there is no StatMan service currently connected to the infrastructure.

#10 - Sep 01, 2017 03:47 PM - Costantino Perciante

Gianpaolo Coro wrote:

One of the big blobs contains a large quantity of objects generated by the old-good Statistical Manager between 2014 and Nov. 2016. These files were written in the ".Application" hidden folder, which is empty for the statistical.manager user. Since this logic was contained in the StatMan service (not in the algorithms execution engine), it is not contained in any way in the DataMiner. Those are old objects that have come from somewhere. Valentina is now checking if they were present also in the Derby DB. Further, there is no StatMan service currently connected to the infrastructure.

I guess you are referring to the largest blobs we found. As far as the other one (i.e. the one in which a stacktrace is reported), I think it would be better to manage that situation in another way (e.g., send it via mail/message), instead of adding the result as property. I hope it can be done

#11 - Oct 16, 2017 11:33 AM - Massimiliano Assante

- Assignee changed from Valentina Marioli to Costantino Perciante

#12 - Oct 18, 2017 01:51 AM - Pasquale Pagano

- Tracker changed from Incident to Task

#13 - Nov 27, 2017 07:32 PM - Andrea Dell'Amico

Some days ago, the DB size increased by almost 10GB in 48 hours. It's been stable again since than, but those increments seem unpredictable and they could be dangerous.

#14 - Feb 21, 2018 04:04 PM - Costantino Perciante

- Assignee changed from Costantino Perciante to Lucio Lelii