

## Infrastructure - Task #8091

### Mysterious sync failure of an eml object from the GOA node

2017-05-05 20:19 - Jing Tao

<b>Status:</b>	In Progress	<b>Start date:</b>	2017-05-05
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Jing Tao	<b>% Done:</b>	30%
<b>Category:</b>	dataone-cn-os-core	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	CCI-2.4.0	<b>Story Points:</b>	
<b>Milestone:</b>	None		
<b>Product Version:</b>	*		
<b>Description</b>			
<p>Matt reported an eml object which is in a package wasn't indexed on production cns but it was indexed on the mn: <a href="https://goa.nceas.ucsb.edu/#view/urn:uuid:3249ada0-afe3-4dd6-875e-0f7928a4c171">https://goa.nceas.ucsb.edu/#view/urn:uuid:3249ada0-afe3-4dd6-875e-0f7928a4c171</a></p> <p>It turned out that the object hadn't been synchronized to cns.</p> <p>I looked at the cn-synchronizaztion.log: [ INFO] 2017-05-05 17:37:18,104 (V2TransferObjectTask:call:234) Task-urn:node:GOA-urn:uuid:3249ada0-afe3-4dd6-875e-0f7928a4c171 - Unlocked Pid. [ERROR] 2017-05-05 17:37:18,104 (V2TransferObjectTask:call:269) Task-urn:node:GOA-urn:uuid:3249ada0-afe3-4dd6-875e-0f7928a4c171 - UnrecoverableException: urn:uuid:3249ada0-afe3-4dd6-875e-0f7928a4c171 cn.createObject failed org.dataone.cn.batch.exceptions.UnrecoverableException: urn:uuid:3249ada0-afe3-4dd6-875e-0f7928a4c171 cn.createObject failed at org.dataone.cn.batch.synchronization.tasks.V2TransferObjectTask.createObject(V2TransferObjectTask.java:808) at org.dataone.cn.batch.synchronization.tasks.V2TransferObjectTask.processNewObject(V2TransferObjectTask.java:493) at org.dataone.cn.batch.synchronization.tasks.V2TransferObjectTask.processTask(V2TransferObjectTask.java:426) at org.dataone.cn.batch.synchronization.tasks.V2TransferObjectTask.call(V2TransferObjectTask.java:211) at org.dataone.cn.batch.synchronization.tasks.V2TransferObjectTask.call(V2TransferObjectTask.java:113) at java.util.concurrent.FutureTask.run(FutureTask.java:266) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142) at java.util.concurrent.ThreadPoolExecutor\$Worker.run(ThreadPoolExecutor.java:617) at java.lang.Thread.run(Thread.java:745) [ WARN] 2017-05-05 17:37:18,106 (SyncFailedTask:submitSynchronizationFailed:116) Task-urn:node:GOA-urn:uuid:3249ada0-afe3-4dd6-875e-0f7928a4c171 - SynchronizationFailed: detail code: 6001 id:urn:uuid:3249ada0-afe3-4dd6-875e-0f7928a4c171 nodeId:urn:node:CNUCSB1 description:Synchronization task of [PID::] urn:uuid:3249ada0-afe3-4dd6-875e-0f7928a4c171 [::PID] failed. Cause: ServiceFailure: Task-urn:node:GOA-urn:uuid:3249ada0-afe3-4dd6-875e-0f7928a4c171 - null [ INFO] 2017-05-05 17:37:18,149 (V2TransferObjectTask:call:294) Task-urn:node:GOA-urn:uuid:3249ada0-afe3-4dd6-875e-0f7928a4c171 - exiting with callState: FAILED</p> <p>I also looked at the metacat.log there is not much information either: metacat 20170505-18:01:17: [ERROR]: D1ResourceHandler: Serializing exception with code 500: null [edu.ucsb.nceas.metacat.restservice.D1ResourceHandler:serializeException:456] metacat 20170505-18:01:17: [WARN]: Could not check docid for revision: urn:uuid:3249ada0-afe3-4dd6-875e-0f7928a4c171 [edu.ucsb.nceas.metacat.EventLog:getReport:291] edu.ucsb.nceas.metacat.McddbDocNotFoundException: the requested docid 'urn:uuid:3249ada0-afe3-4dd6-875e-0f7928a4c171' does not exist</p> <p>metacat 20170505-18:01:17: [ERROR]: D1ResourceHandler: Serializing exception with code 404: No record found for: urn:uuid:3249ada0-afe3-4dd6-875e-0f7928a4c171. [edu.ucsb.nceas.metacat.restservice.D1ResourceHandler:serializeException:456]</p> <p>I asked Peter to uploaded the package into sandbox and the eml object couldn't be sync either. <a href="https://mn-sandbox-ucsb-1.test.dataone.org/knb/d1/mn/v2/query/solr/?q=id:%22urn:uuid:7b45e981-505b-42b1-b598-936da6277ae9%22">https://mn-sandbox-ucsb-1.test.dataone.org/knb/d1/mn/v2/query/solr/?q=id:%22urn:uuid:7b45e981-505b-42b1-b598-936da6277ae9%22</a> <a href="https://cn-sandbox.test.dataone.org/cn/v2/meta/urn:uuid:7b45e981-505b-42b1-b598-936da6277ae9">https://cn-sandbox.test.dataone.org/cn/v2/meta/urn:uuid:7b45e981-505b-42b1-b598-936da6277ae9</a></p>			
<b>Related issues:</b>			
Related to Infrastructure - Bug #8229: cn.creates failing with "413: Request ...		<b>Closed</b>	<b>2017-12-18</b>

---

## History

### #1 - 2017-05-05 20:19 - Jing Tao

- Assignee set to Jing Tao

- Category set to Metacat

### #2 - 2017-05-05 20:22 - Jing Tao

We have a similar sync failure from ebird:

<https://redmine.dataone.org/issues/8054>

### #3 - 2017-05-05 23:48 - Jing Tao

```
[ INFO] 2017-05-05 23:19:12,921 (SyncObjectTask:executeTransferObjectTask:219)
Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 received
[ INFO] 2017-05-05 23:19:12,922 (SyncObjectTask:executeTransferObjectTask:236)
Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 submitted for execution
[DEBUG] 2017-05-05 23:19:12,922 (SyncObjectTask:call:164) ActiveCount: 1 Pool size 1 Max Pool Size 100
[ INFO] 2017-05-05 23:19:12,922 (V2TransferObjectTask:call:202) Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 -
Locking task, attempt 1
[ INFO] 2017-05-05 23:19:12,987 (V2TransferObjectTask:call:207) Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 -
Processing SyncObject
[ INFO] 2017-05-05 23:19:13,509 (V2TransferObjectTask:retrieveMNSystemMetadata:315)
Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 - Retrieved SystemMetadata
Identifier:urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 from node urn:node:mnDemo6 for ObjectInfo Identifier
urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23
[DEBUG] 2017-05-05 23:19:13,509 (V2TransferObjectTask:call:209) Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 MN
system metadata retrieved...
[DEBUG] 2017-05-05 23:19:13,509 (V2TransferObjectTask:processTask:410)
Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 entering processTask...
[DEBUG] 2017-05-05 23:19:13,509 (V2TransferObjectTask:validateSeriesId:579)
Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 entering validateSeriesId...
[DEBUG] 2017-05-05 23:19:13,509 (V2TransferObjectTask:resolve:685) Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23
entering resolve...
[DEBUG] 2017-05-05 23:19:13,824 (V2TransferObjectTask:resolve:695) Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23
PID urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 does not exist on the CN.
[DEBUG] 2017-05-05 23:19:13,824 (V2TransferObjectTask:processNewObject:462)
Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 entering processNewObject...
[DEBUG] 2017-05-05 23:19:13,899 (V2TransferObjectTask:processNewObject:475)
Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 Pid is not reserved by anyone.
[DEBUG] 2017-05-05 23:19:13,899 (V2TransferObjectTask:populateInitialReplicaList:511)
Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 entering populateInitialReplicaList
[DEBUG] 2017-05-05 23:19:13,900 (V2TransferObjectTask:populateInitialReplicaList:526)
Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 Included replica for original MN
[DEBUG] 2017-05-05 23:19:13,900 (V2TransferObjectTask:populateInitialReplicaList:530)
Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 Get Object Format
[DEBUG] 2017-05-05 23:19:13,933 (V2TransferObjectTask:populateInitialReplicaList:542)
Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 Added CN as replica because formatType METADATA is not DATA
[ INFO] 2017-05-05 23:19:13,933 (V2TransferObjectTask:createObject:727)
Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 - Start CreateObject
[ INFO] 2017-05-05 23:19:13,936 (SystemMetadataValidator:schemaValidateSystemMetadata:70) Entering schemaValidateSysMeta method...
[DEBUG] 2017-05-05 23:19:13,937 (V2TransferObjectTask:validateChecksum:835)
Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 entering validateChecksum...
[DEBUG] 2017-05-05 23:19:13,975 (V2TransferObjectTask:createObject:758)
Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 getting ScienceMetadata
[DEBUG] 2017-05-05 23:19:14,016 (V2TransferObjectTask:createObject:795)
Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 Calling CNCreate...
[ INFO] 2017-05-05 23:19:14,055 (V2TransferObjectTask:call:234) Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 -
Unlocked Pid.
[ERROR] 2017-05-05 23:19:14,056 (V2TransferObjectTask:call:269) Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 -
UnrecoverableException: urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 cn.createObject failed
org.dataone.cn.batch.exceptions.UnrecoverableException: urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 cn.createObject failed
at org.dataone.cn.batch.synchronization.tasks.V2TransferObjectTask.createObject(V2TransferObjectTask.java:813)
at org.dataone.cn.batch.synchronization.tasks.V2TransferObjectTask.processNewObject(V2TransferObjectTask.java:491)
at org.dataone.cn.batch.synchronization.tasks.V2TransferObjectTask.processTask(V2TransferObjectTask.java:424)
at org.dataone.cn.batch.synchronization.tasks.V2TransferObjectTask.call(V2TransferObjectTask.java:211)
at org.dataone.cn.batch.synchronization.tasks.V2TransferObjectTask.call(V2TransferObjectTask.java:113)
at java.util.concurrent.FutureTask.run(FutureTask.java:266)
at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142)
at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617)
at java.lang.Thread.run(Thread.java:745)
```

[ WARN] 2017-05-05 23:19:14,057 (SyncFailedTask:submitSynchronizationFailed:116)  
Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 - SynchronizationFailed: detail code: 6001  
id:urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 nodeId:urn:node:cnDevUCSB1 description:Synchronization task of [PID::]  
urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 [::PID] failed. Cause: ServiceFailure:  
Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 - null  
[ INFO] 2017-05-05 23:19:14,075 (V2TransferObjectTask:call:293) Task-urn:node:mnDemo6-urn:uuid:cfe74e5e-d1f9-4fd9-851a-c0d846d15c23 -  
exiting with callState: FAILED

From the above log, the method failed at this line:  
d1Identifier = nodeCommunications.getCnCore().create(session, d1Identifier, sciMetaStream,  
systemMetadata);

But the problem is I can't see something like this at the metacat log:

HTTP Verb: POST

metacat 20170505-22:51:12: [INFO]: D1RestServlet.doPost - HTTP Verb: POST [edu.ucsb.nceas.metacat.restservice.D1RestServlet:doPost:84]

original pathInfo: /object

metacat 20170505-22:51:12: [INFO]: D1HttpRequest.getPathInfo - the original pathInfo: /object

[edu.ucsb.nceas.metacat.restservice.D1HttpRequest:getPathInfo:59]

original requestURI: /metacat/d1/cn/v2/object

So the CN.create wasn't called!

#### #4 - 2017-05-06 03:36 - Jing Tao

- Status changed from New to In Progress

- Category changed from Metacat to dataone-cn-os-core

- % Done changed from 0 to 30

Finally I found the issue - in the apache error log I saw this error:

[Sat May 06 03:23:43.208704 2017] [ssl:error] [pid 31763:tid 140386760611584] [client 128.111.54.78:54623] AH02018: request body exceeds maximum size (131072) for SSL buffer

[Sat May 06 03:23:43.208817 2017] [ssl:error] [pid 31763:tid 140386760611584] [client 128.111.54.78:54623] AH02257: could not buffer message body to allow SSL renegotiation to proceed

When I added the directive on the cn-ssl.conf:

```
SSLRenegBufferSize 100000000
```

After I reloaded the configuration, it worked.

#### #5 - 2017-05-06 03:52 - Jing Tao

- Target version set to CCI-2.4.0

**#6 - 2017-05-09 21:24 - Jing Tao**

Dave just pointed out we may use the http 1.1 header named "Expect". This header essentially says "I've got a huge payload, but before I send it please let me know if you can handle it". This gives the endpoints time to renegotiate the client certificate before the payload is sent.

<http://stackoverflow.com/questions/14281628/ssl-renegotiation-with-client-certificate-causes-server-buffer-overflow>

**#7 - 2017-05-09 21:33 - Jing Tao**

Chris pointed a relative ticket:

<https://redmine.dataone.org/issues/2693>

**#8 - 2017-12-18 23:41 - Rob Nahf**

- Related to Bug #8229: *cn.create failing with "413: Request Entity Too Large" thrown from apache added*