

Infrastructure - Bug #8722

Object in search index but systemmetadata is not available.

2018-10-01 18:38 - Dave Vieglais

Status:	New	Start date:	2018-10-01
Priority:	High	Due date:	
Assignee:	Jing Tao	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:			
Milestone:	None	Story Points:	
Product Version:	*		
Description			
The object: urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 is a CSV file.			
Resolve or getSystemMetadata return a 404 for that object. It is however, available in the search index:			
https://cn.dataone.org/cn/v2/query/solr/?start=0&rows=10&fl=id%3A%22urn%5C%3Aurn%5C%3A4923cca4%5C-c155%5C-4edc%5C-b901%5C-f6e3b4f2e7b2%22			
cn-synchronization.log* reports:			
cn-synchronization.log.32:[INFO] 2018-09-21 15:37:32,418 [ProcessDaemonTask2] (SyncObjectTask:executeTransferObjectTask:293) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 received cn-synchronization.log.32:[INFO] 2018-09-21 15:37:32,418 [ProcessDaemonTask2] (SyncObjectTask:executeTransferObjectTask:310) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 submitted for execution cn-synchronization.log.32:[INFO] 2018-09-21 15:37:32,418 [SynchronizeTask322] (V2TransferObjectTask:call:202) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 - Locking task, attempt 1 cn-synchronization.log.32:[INFO] 2018-09-21 15:37:32,469 [SynchronizeTask322] (V2TransferObjectTask:call:207) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 - Processing SyncObject cn-synchronization.log.32:[INFO] 2018-09-21 15:37:32,575 [SynchronizeTask322] (V2TransferObjectTask:retrieveMNSystemMetadata:317) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 - Retrieved SystemMetadata Identifier:urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 from node urn:node:KNB for ObjectInfo Identifier urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 cn-synchronization.log.32:[INFO] 2018-09-21 15:37:36,547 [SynchronizeTask322] (V2TransferObjectTask:createObject:730) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 - Start CreateObject cn-synchronization.log.32:[INFO] 2018-09-21 15:37:40,388 [SynchronizeTask322] (V2TransferObjectTask:call:234) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 - Unlocked Pid. cn-synchronization.log.32:[ERROR] 2018-09-21 15:37:40,388 [SynchronizeTask322] (V2TransferObjectTask:call:269) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 - UnrecoverableException: urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 cn.createObject failed: The identifier is already in use by an existing object. - InvalidRequest - The identifier is already in use by an existing object. cn-synchronization.log.32:org.dataone.cn.batch.exceptions.UnrecoverableException: urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 cn.createObject failed: The identifier is already in use by an existing object. cn-synchronization.log.32:[WARN] 2018-09-21 15:37:40,389 [SynchronizeTask322] (SyncFailedTask:scheduleSynchronizationFailed:116) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 - SynchronizationFailed: detail code: 6001 id:urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 nodeId:urn:node:CNUCSB1 description:Synchronization task of [PID::] urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 [::PID] failed. Cause: InvalidRequest: The identifier is already in use by an existing object. cn-synchronization.log.32:[INFO] 2018-09-21 15:37:40,460 [SynchronizeTask322] (V2TransferObjectTask:call:294) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 - exiting with call State: FAILED cn-synchronization.log.32:[INFO] 2018-09-21 15:37:40,461 [ProcessDaemonTask2] (SyncObjectTask:re			

```

apFutures:372) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 SyncObjectState: FAILED
cn-synchronization.log.33:[ INFO] 2018-09-21 13:37:17,203 [ProcessDaemonTask2] (SyncObjectTask:executeTransferObjectTask:293) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 received
cn-synchronization.log.33:[ INFO] 2018-09-21 13:37:17,204 [ProcessDaemonTask2] (SyncObjectTask:executeTransferObjectTask:310) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 submitted for execution
cn-synchronization.log.33:[ INFO] 2018-09-21 13:37:17,205 [SynchronizeTask37] (V2TransferObjectTask:call:202) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 - Locking task, attempt 1
cn-synchronization.log.33:[ INFO] 2018-09-21 13:37:17,255 [SynchronizeTask37] (V2TransferObjectTask:call:207) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 - Processing SyncObject
cn-synchronization.log.33:[ INFO] 2018-09-21 13:37:17,462 [SynchronizeTask37] (V2TransferObjectTask:retrieveMNSystemMetadata:317) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 - Retrieved SystemMetadata Identifier:urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 from node urn:node:KNB for ObjectInfo Identifier urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2
cn-synchronization.log.33:[ INFO] 2018-09-21 13:37:21,034 [SynchronizeTask37] (V2TransferObjectTask:createObject:730) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 - Start CreateObject
cn-synchronization.log.33:[ INFO] 2018-09-21 13:37:25,672 [SynchronizeTask37] (V2TransferObjectTask:call:234) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 - Unlocked Pid.
cn-synchronization.log.33:[ERROR] 2018-09-21 13:37:25,673 [SynchronizeTask37] (V2TransferObjectTask:call:269) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 - UnrecoverableException: urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 cn.createObject failed: The identifier is already in use by an existing object. - InvalidRequest - The identifier is already in use by an existing object.
cn-synchronization.log.33:org.dataone.cn.batch.exceptions.UnrecoverableException: urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 cn.createObject failed: The identifier is already in use by an existing object.
cn-synchronization.log.33:[ WARN] 2018-09-21 13:37:25,674 [SynchronizeTask37] (SyncFailedTask:submitSynchronizationFailed:116) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 - SynchronizationFailed: detail code: 6001 id:urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 nodeId:urn:node:CNCSB1 description:Synchronization task of [PID::] urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 [::PID] failed. Cause: InvalidRequest: The identifier is already in use by an existing object.
cn-synchronization.log.33:[ INFO] 2018-09-21 13:37:25,740 [SynchronizeTask37] (V2TransferObjectTask:call:294) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 - exiting with calls state: FAILED
cn-synchronization.log.33:[ INFO] 2018-09-21 13:37:25,740 [ProcessDaemonTask2] (SyncObjectTask:repFutures:372) Task-urn:node:KNB-urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2 SyncObjectState: FAILED

```

Need to determine how an object could be added to the search index, apparently be replicated, but not exist in the system metadata.

History

#1 - 2018-10-01 18:57 - Jing Tao

Hrrm, to me those systemmetadata are available on all three CNs:

<https://cn.dataone.org/cn/v2/meta/urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2>
<https://cn-orc-1.dataone.org/cn/v2/meta/urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2>
<https://cn-unm-1.dataone.org/cn/v2/meta/urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2>

The problem to me is the system metadata only shows four failed replicas. So maybe this is the issue why resolve doesn't work.

#2 - 2018-10-03 17:48 - Jing Tao

Those replicas do exist on mn-ucsb/orc/unm. I used the command to add a replica of KNB:

curl -E /etc/dataone/client/private/urn_node_CNUCSB1.pem -F "replicaMetadata=@replica.xml" -F "serialVersion=73" -X PUT

"<https://cn.dataone.org/cn/v2/replicaMetadata/urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2>"

Use this command to change the status of a replica:

```
curl -E /etc/dataone/client/private/urn_node_CNUCSB1.pem -F "nodeRef=urn:node:mnUCSB1" -F "status=completed" -F "serialVersion=74" -X PUT  
"https://cn.dataone.org/cn/v2/replicaNotifications/urn:uuid:4923cca4-c155-4edc-b901-f6e3b4f2e7b2"
```

Now the resolve command works.