

CN REST - Story #8864

Synchronization does not register authoritative replica entry correctly

2020-06-17 21:49 - Chris Jones

<div>Status:New</div> <div>Priority:Normal</div> <div>Assignee:Jing Tao</div> <div>Category:d1_synchronization</div> <div>Target version:</div> <div>Story Points:</div>	<div>Start date:2020-06-17</div> <div>Due date:</div> <div>% Done:0%</div> <div>Estimated time:0.00 hour</div>
<div>Description</div> <div><p>When objects are synchronized to the CN, the d1_synchronization component will fetch the system metadata for each object and will add a &lt;replica&gt; entry for the origin node (like urn:node:ESS_DIVE, as well as entries for other copies (for instance for science metadata copied to the CN, a &lt;replica&gt;urn:node:CN&lt;/replica&gt; will be added.</p><p>In some instances, the origin replica instance is not added to the replica list.</p><p>This causes downstream problems for the d1_replication component because it relies on the origin node replica entry to be present in order to set up a replication request to a target node. I'm seeing errors like:</p><pre>/var/log/dataone/replicate/cn-replication.log.90:[ERROR] 2020-06-04 05:18:30,179 [pool-15-thread-1] (MNCommunication:requestReplication:34) Could not determine replication source node for replication request for pid: ess-dive-eb6cbb22c605506-20200122T170607966. Replication request failed.</pre></div> <div><p>Looking back in the logs, this is the case for the following objects:</p><pre>ess-dive-3947e68e9825233-20180621T213650539 ess-dive-3b8d9f4513e02f9-20180621T214221437 ess-dive-467a6c3dda4dc88-20180621T211148554 ess-dive-51f345daca126f7-20180328T160350610716 ess-dive-53b37ae5d8c0f51-20200219T211634419654 ess-dive-6b688fab5524c46-20200121T210154766 ess-dive-7a31346c154f02b-20200127T155012488 ess-dive-a1fb05cbd903309-20200130T190835651 ess-dive-b420b097851c716-20180523T161714606 ess-dive-ba81a8a8e0bef31-20180727T200828345 ess-dive-bfaf3d6d6fd038c-20180716T154005175903 ess-dive-c2ef5f3af108c9c-20180621T220020545 ess-dive-eb6cbb22c605506-20200122T170607966 ess-dive-f3238db16593de5-20180621T215956950</pre></div> <div><p>We need to fix this issue in d1_synchronization so replication runs correctly and monthly replica auditing (done by ESS_DIVE) doesn't flag these issues.</p></div>	

History

#1 - 2020-06-19 21:59 - Jing Tao

A replica with the node id - "ess-dive" has been added into the system metadata of those pids.