

Member Nodes - Task #7848

MNDeployment # 3521 (Operational): SEAD Member Node

SEAD in Stage - slow indexing

2016-07-18 22:51 - Laura Moyers

Status:	Closed	Start date:	2016-07-18
Priority:	Normal	Due date:	
Assignee:	Dave Vieglaiss	% Done:	100%
Category:		Estimated time:	0.00 hour
Target version:	Operational		
Story Points:			
Description			
<p>On 7/13/16, Chris, Dave and others discovered a configuration problem in stage where the CNs at UNM and ORC didn't recognize the CN at UCSB:</p> <p>cn-stage-ucsb-1 had a server.name property in metacat set to 'cn-stage.test.dataone.org' instead of cn-stage-ucsb-1.test.dataone.org</p> <p>Chris sorted this out, and indexing of SEAD content began.</p> <p>The current issue is that on Thursday, 7/14/16, only 61 of 208 possible objects had indexed. On Friday, 7/15/16, 61 of 209 possible objects had indexed. On Monday, 7/18/16, 97 objects had indexed of 210 possible objects.</p> <p>Can someone please investigate why this process is going so slowly? This doesn't seem normal.</p>			

History

#1 - 2016-07-21 13:46 - Dave Vieglaiss

- Status changed from New to In Progress
- % Done changed from 0 to 30

Indexing appears to be slow on stage because there are a large number of documents that are not present on the ORC node (which is where indexing is being run). There are 59271 docs in /var/metacat/data on UCSB, but 58786 on ORC

On ORC, cn-index-processor-load.log shows a lot of failed entries:
2016-07-21 13:40:20 new:0, failed: 12571

Switching off indexing on ORC and starting up on UCSB.

On UCSB, cn-index-processor-load.log shows:
2016-07-21 13:43:30 new:43, failed: 2402

Conclusion:

1. There is a lot of content that has not been replicated by metacat
2. There are a lot of index failed errors
3. The slow indexing on ORC is likely a result of the very high number of failed index operations (which get added back onto the queue)

Recommendation:

Recommend a complete purge and re-harvest of the stage environment, adding back one node at a time and verifying sync, replication and indexing tasks proceed as expected.

#2 - 2016-07-25 12:08 - Laura Moyers

As of 7/24/16, SEAD has 214 objects on the MN, 210 have sync'd, and 150 have indexed. It's been at 150 since 7/21/16.

#3 - 2016-10-21 11:20 - Dave Viegla

- *Status changed from In Progress to Closed*

- *% Done changed from 30 to 100*