

Member Nodes - Task #7773

MNDeployment # 3748 (Operational): PPBIO Member Node

Task # 6187 (Closed): PPBIO: Testing

Task # 6244 (Closed): PPBIO: Transition to production

Test content in PPBio in production

2016-05-06 18:12 - Laura Moyers

<b>Status:</b>	Closed	<b>Start date:</b>	2016-05-06
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Mark Servilla	<b>% Done:</b>	100%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	Operational		
<b>Story Points:</b>			
<b>Description</b>			
There is test content in the production PPBio MN. They are re-using their test instance as their production MN, so the test content carried over from stage.			
See list of identifiers (attached) for records generated during testing in stage. These records should be archived before announcing PPBio as a production MN.			

History

#1 - 2016-05-11 22:00 - Mark Servilla

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

#2 - 2016-05-11 22:14 - Mark Servilla

- File PPBIO-test.txt added

There are 104 "test" objects that are in urn:node:PPBIO (see attachment) and, through synchronization, in the CN production environment. Because these objects have a replication policy set to , there are replicas at minimum 2 other MNs.

Dave Vieglais has provided the following discussion and removal instructions:

When archiving an object, all replicas will (should, by design) also be archived by the system. Basically, an object is archived on the MN, the CN is notified and the CN then notifies all replicas that the system metadata has changed. It should only be necessary to set the archive flag on the authoritative MN.

The authoritative MN for an object can be determined by examining the system metadata, the "authoritativeMemberNode" contains the node ID.

The originMemberNode field indicates where the content originally came from. In many cases this will be the same as the authoritativeMemberNode.

All replicas of content have the same system metadata, hence they all have the same replication policy. There is no separate replication policy held by replicas.

To resolve this particular situation with PPBio, the process will be something like:

1. Get a list of test PIDs (origin MN is PPBio and the PID matches a pattern for test identifiers)
2. For each PID found in #1, verify that it is test content
3. Set the archive flag for each test PID by calling mn.archive
4. Verify that the test PIDs no longer appear in searches.

We then need to purge the content using the CN.delete method for each of those PIDs. Technically it is not necessary to do the archive first, but I think in general it is a procedure that should be followed to ensure the content is quickly removed from the search index prior to full deletion.

**#3 - 2016-05-12 14:51 - Mark Servilla**

Tim Vincent (PPBIO MN operator) approved direct access to the urn:node:PPBIO for archiving/removal of inconsequential test objects.

Archival on urn:node:PPBIO for 104 test objects now complete as per instruction #3.

**#4 - 2016-05-12 14:52 - Mark Servilla**

- File *PPBIO-test-archive.txt* added

PPBIO archive log output.

**#5 - 2016-05-12 17:14 - Mark Servilla**

- File *PPBIO-test-delete.txt* added
- translation missing: *en.field\_remaining\_hours* set to 0.0
- % Done changed from 30 to 100
- Status changed from *In Progress* to *Closed*

Deletion of the PPBIO test objects is now complete -- encountered no errors.

Files			
PPBio prod test records from stage.xlsx	10.1 KB	2016-05-06	Laura Moyers
PPBIO-test.txt	3.98 KB	2016-05-11	Mark Servilla
PPBIO-test-archive.txt	31 KB	2016-05-12	Mark Servilla
PPBIO-test-delete.txt	64.3 KB	2016-05-12	Mark Servilla