

Member Nodes - Task #8482

MNDeployment # 8186 (Operational): ESS-DIVE, including CDIAC

Story # 8479 (Closed): ESSDIVE: Testing & Development

ESSDIVE: Test Registration

2018-03-05 15:14 - Amy Forrester

Status:	Closed	Start date:	2018-03-05
Priority:	Normal	Due date:	
Assignee:	Chris Jones	% Done:	100%
Category:		Estimated time:	0.00 hour
Target version:			
Story Points:			
Description			
Custom nodes will start off by first registering in Sandbox. Nodes using an existing MN software will register in Stage.			
<ul style="list-style-type: none">• Node contact subject approved by D1 admin in DataONE LDAP.• Node software configuration - synchronization enabled if applicable.• Node registration document generated and submitted to Sandbox or Stage CN.• D1 Admin approves node registration on the Sandbox or Stage CN server.• Monitor and verify synchronization, indexing, search behavior.• MN approves display of information in test search interface (https://search-sandbox.test.dataone.org/#data or https://search-stage.test.dataone.org/#data)			
A note for custom implemented DataONE services: It's expected that development and testing will be more iterative. When satisfied with results in sandbox, repeat the process by changing target to cn-stage.			
Related issues:			
Related to Infrastructure - Story #8639: Replication performance is too slow ...		New	2018-07-04

History

#1 - 2018-06-15 17:17 - Chris Jones

The urn:node:mnTestESS_DIVE node has been registered in the Stage environment. I have reviewed the node capabilities document, and Val has updated the node name and the node description to appropriate values. Val turned off the MNReplication service since ESS-DIVE won't be a replication target node.

I approved the node in the Stage environment. We noticed some cut/paste issues with apostrophe's in the node description text, and Val called CN.updateNodeCapabilities() with the corrected text, which worked fine (exercising MN to CN authentication).

The CN is harvesting content from the test server now. Synchronization looks to be running smoothly, except for two errors with SVG files that have an incorrect formatId somewhere in the pipeline (could be an MN issue or a CN issue):

```
cjones@cn-stage-ucsb-1:~$ grep ESS_DIVE /var/log/dataone/synchronize/cn-synchronization.log | grep ERROR
[ERROR] 2018-06-14 23:00:39,175 [SynchronizeTask288] (V2TransferObjectTask:populateInitialReplicaList:555) Task-urn:node:mnTestESS_DIVE-ess-dive-e9f4d1f5e8284c5-20180328T194548781 - format NotFound: The format specified by image/svg xml does not exist at this node. - NotFound - The format specified by image/svg xml does not exist at this node.
[ERROR] 2018-06-14 23:00:39,271 [SynchronizeTask288] (V2TransferObjectTask:call:259) Task-urn:node:mnTestESS_DIVE-ess-dive-e9f4d1f5e8284c5-20180328T194548781 - SynchronizationFailed: Synchronization task of [PID::] ess-dive-e9f4d1f5e8284c5-20180328T194548781 [::PID] failed. Cause: NotFound: The format specified by image/svg xml does not exist at this node.
```

We will track down these errors, but all other content has sync'd fine.

#2 - 2018-06-15 17:36 - Chris Jones

Looking at the CN Solr index, only formatType=DATA has indexed correctly on the CN. All other METADATA and RESOURCE types have failed:

```
d1-index-queue=# select formatid, count(formatid) as count from index_task where pid like '%ess-dive%' and status = 'FAILED' group by formatid order by count;
```

formatid	count
http://docs.annotatorjs.org/en/v1.2.x/annotation-format.html	1
FGDC-STD-001.2-1999	23
http://www.openarchives.org/ore/terms	139
eml://ecoinformatics.org/eml-2.1.1	164

Looking at the indexer logs, the indexer is unable to find the science metadata or resource map documents on cn-orc-1 where the indexing is happening. However, the content is on cn-ucsb-1. This indicates a CN to CN replication issue. Looking at the replication log, we in fact do have some errors:

```
metacat 2018-06-14T22:00:02: [INFO]: ReplicationHandler.updateLastCheckTimeForSingleServer - datexml: <error>Metacat received the replication request. However, Metacat can't find the entry of the client certificate or the server parameter on the request url is registered in the xml_replication table. </error>
metacat 2018-06-14T22:00:02: [ERROR]: ReplicationHandler.updateLastCheckTimeForSingleServer - Failed to update last_checked for server cn-stage-ucsb-1.test.dataone.org/metacat/servlet/replication in db because because null
```

So, it looks like there is a certificate issue for CN to CN replication in Stage, which I will track down.

#3 - 2018-06-15 18:00 - Chris Jones

- % Done changed from 0 to 30
- Assignee set to Chris Jones
- Status changed from New to In Progress

#4 - 2018-06-17 02:07 - Chris Jones

The certificates for CN to CN replication in the Stage environment were configured incorrectly, so the CNs weren't replicating science metadata and resource map content among themselves. This is required for indexing to work. I've fixed this in <https://redmine.dataone.org/issues/8617>, and am waiting for all objects to be replicated across the three CNs. Once that is done, I will check back to see how indexing is doing for the ESS-DIVE content.

#5 - 2018-07-03 19:57 - Chris Jones

All ESS-DIVE content that has synchronized has been indexed after the CN to CN replication certs were fixed, so we're good there.

Our next test involved setting replication policies on ESS-DIVE objects so they allow replication, with requested replica numbers of 2 or 3, and with a preferredMemberNode setting of mnStageUCSB2, which emulates replication to the KNB node in production. Val made these changes, which induced the creation of replication tasks for each object on the CN. All of the tasks failed to replicate because of 1) memory issues on the CNs and 2) lack of target MNs to accept the replicas.

I restarted services on the CNs to temporarily fix the memory issues, and re-configured 3 of the 5 replica MNs (upgraded the Metacat installations, re-generated certs that had expired, etc.) I cleared the replication tasks from the replication database (caching the pids so we could re-create the tasks).

I then evicted all of the ESS-DIVE objects from the CN hazelcast cluster, then called hzSystemMetadata.get() for each so that the replication service would create new tasks. It did, and we had mixed results:

Results 2018-07-01 10:52 MDT:

```
select sm.guid, sm.member_node, sm.status
  from smreplicationstatus sm
  inner join systemmetadata sys on sm.guid = sys.guid
 where replication_allowed = true and
       sm.member_node != sys.authoritative_member_node and
       sm.member_node != 'urn:node:cnStage'
 order by sm.guid;
```

guid	member_node	status
ess-dive-003d775b40b2dc3-20180531T000420198	urn:node:mnStageUCSB3	COMPLETED
ess-dive-003d775b40b2dc3-20180531T000420198	urn:node:mnStageUCSB4	QUEUED
ess-dive-003d775b40b2dc3-20180531T000420198	urn:node:mnStagePISCO	QUEUED
ess-dive-003d775b40b2dc3-20180531T000420198	urn:node:mnStageUCSB2	FAILED
ess-dive-003d775b40b2dc3-20180531T000420198	urn:node:mnTestNCEI	FAILED
ess-dive-1026003449059cb-20180607T155243203	urn:node:mnStageUCSB2	COMPLETED
ess-dive-1026003449059cb-20180607T155243203	urn:node:mnStageUCSB3	COMPLETED
ess-dive-1026003449059cb-20180607T155243203	urn:node:mnStagePISCO	FAILED
ess-dive-105dd24c70fa14a-20180613T184401078	urn:node:mnStageUCSB2	COMPLETED
ess-dive-10ea82bba5ec138-20180607T155241093	urn:node:mnStageUCSB3	COMPLETED
ess-dive-10ea82bba5ec138-20180607T155241093	urn:node:mnTestUIC	FAILED
ess-dive-10ea82bba5ec138-20180607T155241093	urn:node:mnStageUCSB2	COMPLETED
ess-dive-15511f3eec91e3b-20180607T155136116	urn:node:mnStagePISCO	FAILED
ess-dive-15511f3eec91e3b-20180607T155136116	urn:node:mnStageUCSB4	COMPLETED
ess-dive-15511f3eec91e3b-20180607T155136116	urn:node:mnStageUCSB2	COMPLETED
ess-dive-1e51ee20381d44d-20180618T220819355	urn:node:mnStageUCSB4	COMPLETED
ess-dive-1e51ee20381d44d-20180618T220819355	urn:node:mnTestUIC	FAILED
ess-dive-1e51ee20381d44d-20180618T220819355	urn:node:mnStageUCSB2	COMPLETED
ess-dive-1e51ee20381d44d-20180618T220819355	urn:node:mnTestNCEI	FAILED
ess-dive-1e51ee20381d44d-20180625T143228115	urn:node:mnStageUCSB2	COMPLETED
ess-dive-1e51ee20381d44d-20180625T143228115	urn:node:mnStageUCSB3	COMPLETED
ess-dive-31ad4d3fe243aa9-20180613T184223144	urn:node:mnStageUCSB2	COMPLETED
ess-dive-3b3d403509db95e-20180330T184741472	urn:node:mnTestUIC	FAILED
ess-dive-3b3d403509db95e-20180330T184741472	urn:node:mnStageUCSB2	COMPLETED
ess-dive-3b3d403509db95e-20180330T184741472	urn:node:mnTestNCEI	FAILED
ess-dive-3b3d403509db95e-20180330T184741472	urn:node:mnStageUCSB3	COMPLETED
ess-dive-422c6d6dd16996e-20180620T153519623	urn:node:mnTestNCEI	FAILED
ess-dive-422c6d6dd16996e-20180620T153519623	urn:node:mnStageUCSB2	COMPLETED
ess-dive-43cadd0f2e0d65d-20180328T124613592452	urn:node:mnStagePISCO	FAILED
ess-dive-43cadd0f2e0d65d-20180328T124613592452	urn:node:mnStageUCSB2	COMPLETED
ess-dive-43cadd0f2e0d65d-20180328T124613592452	urn:node:mnTestNCEI	FAILED
ess-dive-43cadd0f2e0d65d-20180328T124613592452	urn:node:mnStageUCSB3	QUEUED
ess-dive-44e81bb265dc897-20180620T153525681	urn:node:mnStageUCSB2	FAILED
ess-dive-44e81bb265dc897-20180620T153525681	urn:node:mnStageUCSB4	COMPLETED
ess-dive-44e81bb265dc897-20180620T153525681	urn:node:mnTestNCEI	FAILED
ess-dive-60a2e01b3dbee05-20180613T184358767	urn:node:mnStageUCSB2	COMPLETED
ess-dive-641402af0c3deae-20180328T194004284	urn:node:mnStagePISCO	FAILED
ess-dive-641402af0c3deae-20180328T194004284	urn:node:mnTestUIC	FAILED
ess-dive-641402af0c3deae-20180328T194004284	urn:node:mnStageUCSB2	FAILED
ess-dive-6db0e8dfd3191cd-20180328T124605561744	urn:node:mnTestUIC	FAILED
ess-dive-6db0e8dfd3191cd-20180328T124605561744	urn:node:mnStagePISCO	FAILED
ess-dive-6db0e8dfd3191cd-20180328T124605561744	urn:node:mnTestNCEI	FAILED
ess-dive-6db0e8dfd3191cd-20180328T124605561744	urn:node:mnStageUCSB2	FAILED
ess-dive-8b533c8de9ce3ea-20180531T000423664	urn:node:mnStageUCSB2	FAILED
ess-dive-968f4b8b5eb06ca-20180330T184732949	urn:node:mnTestUIC	FAILED
ess-dive-968f4b8b5eb06ca-20180330T184732949	urn:node:mnStageUCSB2	FAILED
ess-dive-968f4b8b5eb06ca-20180330T184732949	urn:node:mnStagePISCO	FAILED

```

ess-dive-b3ef2dfcbb9c881-20180618T220826574 | urn:node:mnTestUIC | FAILED
ess-dive-b3ef2dfcbb9c881-20180618T220826574 | urn:node:mnStagePISCO | FAILED
ess-dive-b3ef2dfcbb9c881-20180618T220826574 | urn:node:mnStageUCSB2 | FAILED
ess-dive-b8e3670bcf05b7f-20180625T143234735 | urn:node:mnStageUCSB2 | FAILED
ess-dive-be67a7e3e2f54a7-20180328T202113712 | urn:node:mnTestUIC | FAILED
ess-dive-be67a7e3e2f54a7-20180328T202113712 | urn:node:mnStageUCSB2 | FAILED
ess-dive-be67a7e3e2f54a7-20180328T202113712 | urn:node:mnStagePISCO | FAILED
ess-dive-c4175d8a6998e50-20180328T193626054 | urn:node:mnStagePISCO | FAILED
ess-dive-c4175d8a6998e50-20180328T193626054 | urn:node:mnTestUIC | FAILED
ess-dive-c4175d8a6998e50-20180328T193626054 | urn:node:mnStageUCSB2 | FAILED
ess-dive-c4175d8a6998e50-20180328T193626054 | urn:node:mnTestNCEI | FAILED
ess-dive-fba3f9dc07486f6-20180328T202708237 | urn:node:mnStageUCSB3 | FAILED
ess-dive-fba3f9dc07486f6-20180328T202708237 | urn:node:mnTestUIC | FAILED
ess-dive-fba3f9dc07486f6-20180328T202708237 | urn:node:mnStageUCSB2 | FAILED
ess-dive-fba3f9dc07486f6-20180328T202708237 | urn:node:mnTestNCEI | FAILED
ess-dive-ff3548ec10ba92a-20180330T192903805 | urn:node:mnTestNCEI | FAILED
ess-dive-ff3548ec10ba92a-20180330T192903805 | urn:node:mnStageUCSB2 | FAILED
(64 rows)

```

To summarize:

```

mnStageUCSB2: 11 COMPLETED, 12 FAILED
mnStageUCSB3:  5 COMPLETED,  1 FAILED
mnStageUCSB3:  3 COMPLETED,  0 FAILED
  mnTestNCEI:  0 COMPLETED, 10 FAILED
  mnTestUIC:  0 COMPLETED, 10 FAILED
  mnTestPISCO: 0 COMPLETED,  9 FAILED

```

The failures replicating to mnTestUIC, mnTestPISCO, and mnTestNCEI can be explained by issues on the target MNs. The failures on mnStageUCSB{2-4} need to be investigated, since we had many successes on those newly-reconfigured target nodes. We need to determine if the issues originate on the source MN (ESS-DIVE), the CN, or the target MN. I will be working on that, along with a number of replication service issues that have arisen since Skye refactored the code to use a SQL-based task queue.

#6 - 2018-07-03 23:11 - Chris Jones

After looking more closely at some of the FAILED rows above, I'm seeing that the CN has different results now:

Results 2018-07-03 05:07 PM MDT:

```
select sm.guid, sm.member_node, sm.status
  from smreplicationstatus sm
  inner join systemmetadata sys on sm.guid = sys.guid
  where replication_allowed = true and
        sm.member_node != sys.authoritative_member_node and
        sm.member_node != 'urn:node:cnStage' and sys.guid like 'ess-dive%'
  order by sm.guid;
```

guid	member_node	status
ess-dive-003d775b40b2dc3-20180531T000420198	urn:node:mnStagePISCO	FAILED
ess-dive-003d775b40b2dc3-20180531T000420198	urn:node:mnStageUCSB3	COMPLETED
ess-dive-003d775b40b2dc3-20180531T000420198	urn:node:mnStageUCSB4	COMPLETED
ess-dive-003d775b40b2dc3-20180531T000420198	urn:node:mnStageUCSB2	COMPLETED
ess-dive-003d775b40b2dc3-20180531T000420198	urn:node:mnTestNCEI	FAILED
ess-dive-1026003449059cb-20180607T155243203	urn:node:mnStagePISCO	FAILED
ess-dive-1026003449059cb-20180607T155243203	urn:node:mnStageUCSB2	COMPLETED
ess-dive-1026003449059cb-20180607T155243203	urn:node:mnStageUCSB3	COMPLETED
ess-dive-105dd24c70fa14a-20180613T184401078	urn:node:mnStageUCSB2	COMPLETED
ess-dive-10ea82bba5ec138-20180607T155241093	urn:node:mnTestUIC	FAILED
ess-dive-10ea82bba5ec138-20180607T155241093	urn:node:mnStageUCSB2	COMPLETED
ess-dive-10ea82bba5ec138-20180607T155241093	urn:node:mnStageUCSB3	COMPLETED
ess-dive-15511f3eec91e3b-20180607T155136116	urn:node:mnStagePISCO	FAILED
ess-dive-15511f3eec91e3b-20180607T155136116	urn:node:mnStageUCSB4	COMPLETED
ess-dive-15511f3eec91e3b-20180607T155136116	urn:node:mnStageUCSB2	COMPLETED
ess-dive-1e51ee20381d44d-20180618T220819355	urn:node:mnStageUCSB4	COMPLETED
ess-dive-1e51ee20381d44d-20180618T220819355	urn:node:mnTestUIC	FAILED
ess-dive-1e51ee20381d44d-20180618T220819355	urn:node:mnStageUCSB2	COMPLETED
ess-dive-1e51ee20381d44d-20180618T220819355	urn:node:mnTestNCEI	FAILED
ess-dive-1e51ee20381d44d-20180625T143228115	urn:node:mnStageUCSB2	COMPLETED
ess-dive-1e51ee20381d44d-20180625T143228115	urn:node:mnStageUCSB3	COMPLETED
ess-dive-31ad4d3fe243aa9-20180613T184223144	urn:node:mnStageUCSB2	COMPLETED
ess-dive-3b3d403509db95e-20180330T184741472	urn:node:mnStageUCSB2	COMPLETED
ess-dive-3b3d403509db95e-20180330T184741472	urn:node:mnTestUIC	FAILED
ess-dive-3b3d403509db95e-20180330T184741472	urn:node:mnTestNCEI	FAILED
ess-dive-3b3d403509db95e-20180330T184741472	urn:node:mnStageUCSB3	COMPLETED
ess-dive-422c6d6dd16996e-20180620T153519623	urn:node:mnTestNCEI	FAILED
ess-dive-422c6d6dd16996e-20180620T153519623	urn:node:mnStageUCSB2	COMPLETED
ess-dive-43cadd0f2e0d65d-20180328T124613592452	urn:node:mnStageUCSB2	COMPLETED
ess-dive-43cadd0f2e0d65d-20180328T124613592452	urn:node:mnTestNCEI	FAILED
ess-dive-43cadd0f2e0d65d-20180328T124613592452	urn:node:mnStageUCSB3	COMPLETED
ess-dive-43cadd0f2e0d65d-20180328T124613592452	urn:node:mnStagePISCO	FAILED
ess-dive-44e81bb265dc897-20180620T153525681	urn:node:mnStageUCSB2	FAILED
ess-dive-44e81bb265dc897-20180620T153525681	urn:node:mnStageUCSB4	COMPLETED
ess-dive-44e81bb265dc897-20180620T153525681	urn:node:mnTestNCEI	FAILED
ess-dive-60a2e01b3dbe05-20180613T184358767	urn:node:mnStageUCSB2	COMPLETED
ess-dive-641402af0c3deae-20180328T194004284	urn:node:mnStageUCSB2	FAILED
ess-dive-641402af0c3deae-20180328T194004284	urn:node:mnTestUIC	FAILED
ess-dive-641402af0c3deae-20180328T194004284	urn:node:mnStagePISCO	FAILED
ess-dive-6db0e8dfd3191cd-20180328T124605561744	urn:node:mnTestUIC	FAILED
ess-dive-6db0e8dfd3191cd-20180328T124605561744	urn:node:mnStagePISCO	FAILED
ess-dive-6db0e8dfd3191cd-20180328T124605561744	urn:node:mnTestNCEI	FAILED
ess-dive-6db0e8dfd3191cd-20180328T124605561744	urn:node:mnStageUCSB2	FAILED
ess-dive-8b533c8de9ce3ea-20180531T000423664	urn:node:mnStageUCSB2	FAILED
ess-dive-968f4b8b5eb06ca-20180330T184732949	urn:node:mnTestUIC	FAILED
ess-dive-968f4b8b5eb06ca-20180330T184732949	urn:node:mnStageUCSB2	FAILED
ess-dive-968f4b8b5eb06ca-20180330T184732949	urn:node:mnStagePISCO	FAILED
ess-dive-b3ef2dfcbb9c881-20180618T220826574	urn:node:mnTestUIC	FAILED
ess-dive-b3ef2dfcbb9c881-20180618T220826574	urn:node:mnStagePISCO	FAILED
ess-dive-b3ef2dfcbb9c881-20180618T220826574	urn:node:mnStageUCSB2	FAILED
ess-dive-b8e3670bcf05b7f-20180625T143234735	urn:node:mnStageUCSB2	FAILED
ess-dive-be67a7e3e2f54a7-20180328T202113712	urn:node:mnTestUIC	FAILED
ess-dive-be67a7e3e2f54a7-20180328T202113712	urn:node:mnStageUCSB2	FAILED
ess-dive-be67a7e3e2f54a7-20180328T202113712	urn:node:mnStagePISCO	FAILED
ess-dive-c4175d8a6998e50-20180328T193626054	urn:node:mnStagePISCO	FAILED
ess-dive-c4175d8a6998e50-20180328T193626054	urn:node:mnTestUIC	FAILED
ess-dive-c4175d8a6998e50-20180328T193626054	urn:node:mnStageUCSB2	FAILED
ess-dive-c4175d8a6998e50-20180328T193626054	urn:node:mnTestNCEI	FAILED
ess-dive-fba3f9dc07486f6-20180328T202708237	urn:node:mnStageUCSB3	FAILED
ess-dive-fba3f9dc07486f6-20180328T202708237	urn:node:mnTestUIC	FAILED

```
ess-dive-fba3f9dc07486f6-20180328T202708237 | urn:node:mnStageUCSB2 | FAILED
ess-dive-fba3f9dc07486f6-20180328T202708237 | urn:node:mnTestNCEI | FAILED
ess-dive-ff3548ec10ba92a-20180330T192903805 | urn:node:mnTestNCEI | FAILED
ess-dive-ff3548ec10ba92a-20180330T192903805 | urn:node:mnStageUCSB2 | FAILED
```

So more replica statuses are COMPLETED for the mnStageUCSB{2,3,4} nodes. I'll look into the mismatch.

#7 - 2018-07-04 11:18 - Dave Vieglais

- *Related to Story #8639: Replication performance is too slow to service demand added*

#8 - 2018-07-04 15:39 - Chris Jones

It was only one new updated status per mnStage{2,3,4} machine, so I was probably just too quick to report back:

```
mnStageUCSB2: 12 COMPLETED, 11 FAILED
mnStageUCSB3: 6 COMPLETED, 1 FAILED
mnStageUCSB3: 4 COMPLETED, 0 FAILED
```

#9 - 2018-08-06 22:52 - Chris Jones

- *% Done changed from 30 to 100*

- *Status changed from In Progress to Closed*

The ESS-DIVE node was successfully registered in STAGE, and all content synchronized as expected. We had some glitches in replication of content, but all issues revolved around CN configuration and efficiency or the lack of MN targets. Eventually all objects set to replicate did so successfully.