

Infrastructure - Bug #6773

PID Discrepancy on production CNs

2015-01-26 17:52 - Dave Vieglais

Status: Closed	Start date: 2015-01-26
Priority: Urgent	Due date:
Assignee: Robert Waltz	% Done: 100%
Category: Environment.Production	Estimated time: 0.00 hour
Target version:	Story Points:
Milestone: None	
Product Version: *	

Description

It looks like a discrepancy across the production CNs started on 2015-01-23:

```
|. Time          |. UCSB |. ORC |. UNM |
| 2015-01-26T05:15:25Z | 577915 | 577915 | 577831 |
| 2015-01-25T05:15:23Z | 577912 | 577912 | 577828 |
| 2015-01-24T05:15:31Z | 577652 | 577652 | 577567 |
| 2015-01-23T05:15:28Z | 577410 | 577410 | 577403 |
| 2015-01-22T05:15:26Z | 577202 | 577202 | 577202 |
| 2015-01-21T05:15:27Z | 576946 | 576946 | 576946 |
```

This list of identifiers is available at:

https://monitor.dataone.org/_production_history/data/CN-PID-comparison.yaml

Related issues:

Related to Infrastructure - Bug #6779: Metacat "hazelcast synchronization" do...	Closed	2015-01-27
----------------------------------------------------------------------------------	---------------	-------------------

History

#1 - 2015-01-26 19:32 - Robert Waltz

Problem report:

(12:53:47) vieglais: hi Robert - looks like CN-UNM-1 might of dropped out of the cluster

(13:25:21) robert: Dave, CN-UNM-1 periodically has dropped out of the cluster over the past month. It is usually only a couple of minutes, and then it rejoins. Given its regularity since December and the short timeframe that it is out, I suppose that it is a regular process happening on UNM networks. It happens at least once a week over the past month.

(13:26:38) robert: Dates are - Jan 5,6,13,14, 21,23. I predict by the end of this week it will happen twice.

(13:26:52) vieglais: have you noticed if it is something like a complete disconnect or any particular service?

(13:28:05) robert: I haven't checked LDAP but Mark Servilla did note on one of the days that LDAP was messed up, but he reported that it repaired itself.

(13:28:23) vieglais: previous events, at least over the last month, didn't cause a consistency issue - perhaps just coincidence?

(13:30:45) robert: I don't know. We do have a timeframe of when the inconsistency began. We regained consistency last December, 17th.

(13:31:22) vieglais: looking at check_mk - ldap between unum and ucsb was up and down almost once per minute

(13:31:32) robert: when?

(13:31:35) vieglais: on the 23rd

(13:31:49) vieglais: last event was 2015-01-23 15:24:05

(13:32:52) vieglais: the 22nd was mostly ok

(13:33:11) vieglais: the 21st had similar LDAP dropouts to the 23rd

(13:33:26) robert: that is a finer timeframe than hazelcast checks, by default, hazelcast only checks every 5 minutes for a response to its Master node confirmation request.

(13:45:36) robert: Max timeout of heartbeat in seconds for a node to assume it is dead: 600 (default)

(13:46:36) robert: So, if there are networking problems, and the disruption is less than 5 minutes and the disrupted node rejoins the cluster, we will not see a splunk error report.

(13:47:22) robert: check_mk and LDAP may give us a better indication of the frequency it is really happening.

(13:50:02) vieglais: where does ldap log to on the cns?

#2 - 2015-01-27 07:49 - Robert Waltz

cn-unm-1 restoration of consistency actions:

I. Check & Restore consistency of LDAP on cn-unm-1:

A. On CN-UNM-1

1. Check consistency LDAP

a. Diff Node Lists - If node lists are the same, then compare **fill in the blank** variable on all the Node entries.

b. If results indicate consistency, goto II.

2. Restore consistency LDAP

a. Stop LDAP/restart LDAP - the ldap process can be disconnected and reconnected on the CN components with out needing to stop/restart dependent processes (cn.war and process daemons)

b. Repeat 1.a, if 1.a does not pass, then scratch your head. It might be an issue on other machines, it maybe that the network is still split, it maybe that ldap is taking longer to synchronize than normal (so just sit still and repeat check in 10 minutes)

II. Check & Restore consistency of Metacat on cn-unm-1:

A. On CN-UNM-1

1. Check consistency Metacat

a. Compare number counts on monitor.dataone.org?

b. Perform the following command for each of the CNs:

<https://cn-unm-1.dataone.org/cn/v1/object?start=0&count=0>

<https://cn-ucsb-1.dataone.org/cn/v1/object?start=0&count=0>

<https://cn-orc-1.dataone.org/cn/v1/object?start=0&count=0>

c. If the counts are equal, then goto end. If this is your first time running these commands, What made you check metacat consistency? Please make a note. maybe only a problem in ldap...

2. Restore consistency Metacat

a. Login to the administrative interface of metacat - <https://cn-unm-1.dataone.org/metacat/admin>

b. Click the button named 'Reconfigure Now' beside the label 'Replication Configuration'

c. Perform the 'Hazelcast Synchronization' operation.

i. Begin by tailing the log in /var/metacat/logs/metacat.log (tail -100f /var/metacat/logs/metacat.log)

ii. Click the button named 'Resynch' under the label 'Hazelcast Synchronization'

iii. Note in the log files, the beginning of the synchronization process! (remember that logs rotate, so re-issue the tail command periodically)

metacat 20150126-20:55:51: [WARN]: Local SystemMetadata pid count: 577993

metacat 20150126-20:57:49: [WARN]: processedCount (identifiers from iterator): 578077

There are a lot of these messages! metacat 20150126-20:58:04: [ERROR]: Error looking up missing system metadata for pid:

<http://dx.doi.org/10.5061/dryad.ft48k/3/bitstream> [edu.ucsb.nceas.metacat.dataone.hazelcast.HazelcastService]

is this right? need to check the metacat code to ensure it is nothing fatal.

OK, there is no official end to the process in the log files. It just kind of stops logging. So, lets gamble that everything is complete! ...NOOO! ... Goto

II.A.1.a and start over

```
d. Perform the 'Replicate Now' operation
  i. Begin by tailing the log in /var/metacat/logs/replicate.log (tail -100f /var/metacat/logs/replicate.log)
g)
  ii. Click the button named 'Get All' under the label 'Replicate Now'
  iii. Messages to expect in replicate.log
      Start of the Metacat replication process
      metacat 2015-01-26T14:48:47: [INFO]: ReplicationService.handleForceReplicateDataFileRequest - Force replication request from: cn-ucsb-1.dataone.org/metacat/servlet/replication
      End of the Metacat replication process
      metacat 2015-01-27T07:34:22: [WARN]: ForceReplicationHandler.run - exiting ForceReplicationHandler Thread
```

#3 - 2015-01-27 14:40 - Dave Vieglais

- Status changed from New to Closed

Consistency checks out ok.

Added wiki page describing the process: <https://redmine.dataone.org/projects/d1-servers/wiki/Recipes>

#4 - 2015-01-27 14:44 - Dave Vieglais

- Related to Bug #6779: Metacat "hazelcast synchronization" does not emit a log message indicating the process has completed added