

Infrastructure - Task #7612

Story # 7611 (New): ORNLDAAC fails during log aggregation

Contact ORNLDAAC

2016-01-26 16:39 - Robert Waltz

Status:	In Progress	Start date:	2016-01-26
Priority:	Normal	Due date:	
Assignee:	Laura Moyers	% Done:	30%
Category:	Environment.Production	Estimated time:	0.00 hour
Target version:	Maintenance Backlog	Story Points:	
Milestone:	None		
Product Version:	*		

Description

ORNLDAAC administrative contact needs to be informed that the MemberNode is failing during a REST request call to its logging endpoint. After ORNLDAAC has fixed the issue, then D1 will need to evaluate log aggregation responses before resetting log aggregation to harvest node.

History

#1 - 2016-01-26 20:08 - Laura Moyers

- Status changed from New to In Progress

- % Done changed from 0 to 30

Contacted Bob Cook and Giri Palanisamy; awaiting response.

#2 - 2016-02-01 17:44 - Robert Waltz

Not sure what the goal is here, but it looks like D1 is trying pull with a count greater than the available results.

For the specified dates, the count is 22. So, 1,000 returns an error.

From Bob Cook

#3 - 2016-02-01 21:50 - Dave Vieglais

If the requested number of objects is larger than the number available, then the response should include those available.

#4 - 2016-02-11 17:36 - Robert Waltz

I tested another Mercury instance with a different query to determine if we receive the same type of error:

With curl I can query "

http://mercury-ops2.ornl.gov/RGD_MN/mn/v1/log?fromDate=2014-10-21T11:54:00.000%2B00:00&toDate=2016-02-11T00:00:00.000%2B00:00&start=0&count=0"

and receive the results

With curl I can query "

http://mercury-ops2.ornl.gov/RGD_MN/mn/v1/log?fromDate=2014-10-21T11:54:00.000%2B00:00&toDate=2016-02-11T00:00:00.000%2B00:00&start=0&count=1000"

and receive 393 records back without error.

#5 - 2016-02-11 17:41 - Robert Waltz

As Dave noted, the count parameter shows "The maximum number of log records that should be returned in the response. The Member Node may return fewer and the caller should check the total in the response to determine if further pages may be retrieved."

To answer Bob's question, the goal is to harvest log records from ORNLDAAC through the DataONE API. It appears that ORNLDAAC is incorrectly interpreting count as 'an exact number of log records that should be returned', if indeed the error is due to the requested count being larger than the actual number of records to be returned.

#6 - 2016-02-11 18:02 - Robert Waltz

Using curl, here are other scenarios that are similar to the original error:

With curl I can query ORNLDAAC for all of its records:

<http://mercury-ops2.ornl.gov/ornldaac/mn/v1/log?start=0&count=0>

ORNLDAAC returns:

If I just ask for a small subset at the end of that query, using start of 8800 and count of 1000:

<http://mercury-ops2.ornl.gov/ornldaac/mn/v1/log?start=8800&count=1000>

I do receive 32 records back.

Similarly, if I add dates to the above queries:

"

<http://mercury-ops2.ornl.gov/ornldaac/mn/v1/log?fromDate=1900-01-01T00:00:00.000%2B00:00&toDate=2016-02-11T00:00:00.000%2B00:00&start=0&count=0>"

"

<http://mercury-ops2.ornl.gov/ornldaac/mn/v1/log?fromDate=1900-01-01T00:00:00.000%2B00:00&toDate=2016-02-11T00:00:00.000%2B00:00&start=8800&count=1000>"

I do receive 32 records back.