

Infrastructure - Bug #2356

No Nodeld for the virtual CN represented by the RR DNS entry for CNs

2012-02-21 17:19 - Dave Vieglaiss

Status:	Closed	Start date:	2012-02-21
Priority:	High	Due date:	
Assignee:	Robert Waltz	% Done:	100%
Category:	d1_synchronization	Estimated time:	0.00 hour
Target version:	Sprint-2012.09-Block.2.1	Story Points:	
Milestone:	CCI-1.0.0		
Product Version:	*		
Description			
<p>During synchronization a entry is made for science metadata content indicating that the content has been replicated to a Coordinating Node. However, since the content is actually replicated to all CNs within the environment only a single entry is made. The nodeld for this replica is undefined since it should refer to the "virtual" CN indicating by the RR DNS entry for the environment so that the object can be retrieved from any CN. However, there is no nodelID that corresponds to the RR DNS entry.</p> <p>See #2355 for example.</p> <p>Possible resolutions:</p> <ol style="list-style-type: none">1. Create a new nodelID that refers to the RR DNS entry2. Use a "magic" value for nodelID that is interpreted by the resolve() method to mean any CN3. Use the nodelID of the CN where synchronization was made4. Add entries for each CN holding the object. <p>Option 4 is perhaps most appropriate, and the entry should be made when the object appears on the respective CN.</p> <p>Option 3 is the simplest fix that will enable resolve() to function correctly.</p> <p>Option 2 is a hack, and may have consequences for replica management.</p> <p>Option 1 is probably asking for trouble.</p> <p>Suggestion is to modify synchronization to use the current nodeld (i.e. option 3) until a better solution is determined.</p>			
Related issues:			
Related to Infrastructure - Bug #2355: Incorrect node id added to replica ent...		Closed	2012-02-21

History

#1 - 2012-02-24 17:57 - Robert Waltz

- Status changed from New to In Progress

#2 - 2012-02-27 18:20 - Dave Vieglaiss

- Target version changed from Sprint-2012.07-Block.1.4 to Sprint-2012.09-Block.2.1

- Position set to 37

#3 - 2012-02-29 06:13 - Robert Waltz

- Status changed from In Progress to Closed

Actually Option [#1](#) is what we decided to do a few months ago and what has been implemented. I never populated the nodelist until now because we didn't have round robin.

We need to be able to access the RR DNS entry from the libclient packages, so we have to have it in the node list regardless.

We can change the entry pretty easily through the postinst scripts if we need to implement option 3.

Option 4 would be very easy to implement as well and would need no configuration changes. The issue with [#4](#) is when we add a new CN. We would

then have to modify all systemMetadata to include the new CN as an entry.

Having the round robin entry as the node ensures that a client can retrieve the data from any CN.