

Infrastructure - Task #4106

Configure PostgreSQL memory on the CNs to accommodate large getLogRecords() calls

2013-10-18 21:03 - Chris Jones

<b>Status:</b>	Closed	<b>Start date:</b>	2013-10-18
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Chris Brumgard	<b>% Done:</b>	100%
<b>Category:</b>	d1_cn_buildout	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	2013.44-Block.6.1	<b>Story Points:</b>	
<b>Milestone:</b>	None		
<b>Product Version:</b>	*		
<b>Description</b>			
<p>For Metacat-based MN and CN installations with millions of log records, calls to getLogRecords() take over 5 minutes to execute, and the limiting factor seems to be the SQL query. Ben has patched Metacat with two new indices in the database, which enables faster queries against the access_log table. However, this needs to be coupled with changes to the PostgreSQL memory configuration.</p> <p>For a 32GB VM, we set the effective_cache_size to 16GB so the query planner would prioritize using indices and taking advantage of available RAM. We also set work_mem = 100MB, with max_connections at 100. This would result in a potential use of 10GB of RAM if all 100 connections were executing large or complex sort operations.</p> <p>For the CNs, change these settings in dataone-cn-os-core postinst script to appropriate levels. For instance, if max_connections is 200, perhaps use work_mem = 50MB on a 32GB VM. The recommended settings for effective_cache_size is 1/2 to 3/4 of total RAM on the VM.</p> <p>These changes should coincide with an upgrade to Metacat 2.2.2 on the CNs.</p> <p>See <a href="http://wiki.postgresql.org/wiki/Tuning_Your_PostgreSQL_Server">http://wiki.postgresql.org/wiki/Tuning_Your_PostgreSQL_Server</a></p>			
<b>Related issues:</b>			
Related to Infrastructure - Story #4193: Release CCI 1.2.4 Features		Closed	2013-12-11 2013-12-14

History

#1 - 2013-10-22 00:43 - Chris Brumgard

Chris,

This should be easy enough to do.

1) When is the upgrade of Metacat to 2.2.2 planned?

2) Which set of CN will this affect? If this is the prod boxes, I only have sudo privileges on the orc node. I'm not suppose to have sudo privileges on any of the prod boxes after our security discussion a few weeks ago.

#2 - 2013-11-01 14:21 - Chris Brumgard

For configuring the work memory, should work memory be always set to 10GB / max\_connections or 1/3 RAM / max\_connections.

#3 - 2013-11-01 14:42 - Chris Jones

Chris - the point is to allow plenty of work\_mem for sorting operations, etc., and so in lieu of the VMs total RAM, 25% to 33%, divided by max\_connections, would be a decent work\_mem setting. However, for VMs with little memory (say 4GB), I'd be more conservative and use 15-20%.

#4 - 2013-11-05 21:02 - Chris Brumgard

- Status changed from New to Closed

- % Done changed from 0 to 100
- translation missing: en.field\_remaining\_hours set to 0.0