

CMS dCache - Bug #8253

Reduce dcache maximum memory to numbers within the server available memory

04/06/2015 02:45 PM - Natalia Ratnikova

Status:	Resolved	Start date:	04/06/2015
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.50 hour
component:		Stakeholders:	
First Occurred:		Co-Assignees:	
Occurs In:			

Description

Based on request originated from INC000000523602 and RITM0188297 in Service-now.

There are three nodes in this condition: 8GB of physical memory and three pools each configured with 5Gb memory reserved for dCache.

All are OOW and belong to the tape instance :

cmsstor108
cmsstor110
cmsstor112

Gerard's proposal is included in [1] below.

Solution based on input from dCache support meeting are summarized by Natalia in [2].

[1]

2015-03-26 05:03:19 Gerard Bernabeu Altayo - Changed: Additional comments

[Additional comments] I think this does not require immediate intervention, the system may be able to recover on its own, if it turns to become an issue I suggest restarting dcache, if it does not help much reboot the box.

What I really want to do is:

1. Modify modules/dcache/templates/layout-pool.erb so that the following values:

```
dcache.java.memory.heap=5096m  
dcache.java.memory.direct=5096m
```

Are actually $\$::memorysize / \${poolcount}$.

Where memorysize is a system fact and I should be able to get the # of pools from somewhere.

This will imply rebooting all pools, so I need to do/test/plan this a bit.

Gerard

[2]

2015-03-26 14:47:07 Natalia Ratnikova - Changed: Additional comments

[Additional comments] Hi Gerard,

I discussed this today at dCache support meeting. The developers confirmed that over-committed dCache memory settings may cause problems and should be fixed.

Normally there is no need to go over 6-8 GB memory requirement per pool, unless there is a very large number of small files. This agrees with recommendations Dmitry gave us earlier.

Going over 32Gb may be harmful due to a larger address length. Xavier reported that they setup 2GB per pool at KIT for CMS, and did not experience any problems.

Given this, I do not see a need to customize the heap size per each pool. All we need to do is to make sure that we do not over-commit.

If the limit is hit, dCache will restart the pool, produce the heap dump and log OOM exception. We have zabbix monitoring for this and will know if this happens.

As concerns cmsstor110 node in particular, I found it in cms-retired list. So probably the easiest thing is to just decommission this node.

I checked with the ECF-CIS group for a full list of cmsstor* nodes to be retired.

They will send their weekly reports on the nodes "to-be-retired" to dcs0@fnal.gov (approved by Glenn).

- Natalia.

History

#1 - 04/06/2015 02:51 PM - Natalia Ratnikova

[Added Gerard on the watch list]

Since we decided to keep cmsstor110 in production, and cmsstor108 cmsstor112 have the same "mishap", we may want to have a custom config for these three nodes only: 2Gb per relatively small pool should be enough.

#2 - 05/14/2015 02:14 PM - Natalia Ratnikova

Three nodes in Tape instance that had too small memory have been retired by Chih-Hao.

#3 - 05/14/2015 02:15 PM - Natalia Ratnikova

- *Status changed from New to Resolved*