

Instability of node 3 and 4 of CMSR database affecting online to offline replication

Description

On Friday 20th August in the morning nodes 3 and 4 of the CMSR production database rebooted few times causing hand of Streams replication.

Impact

- The reboots affected directly CMS PhEDEx and CMS Dataset Bookkeeping applications which are main users of those nodes. The applications lost connections and transactions in progress. They could reconnect to the database immediately thanks to service failover mechanism.
- Streams replication hangs could potentially affect re-processing activities which are relying on streaming of condition data.

Time line of the incident

All events described below happened on 20.08.2010.

- 9:12 - cmsr3 reboots -> node rebooted by cluster, services relocated to remaining nodes.
- 9:30 - high load on cmsr4 due to 2 user processes consuming all the memory of the server. The server was completely overloaded and affected new connections to the cmsr cluster.
- 9:40 - cmsr4 cleared -> the issue of high load of cmsr4 was cleared by killing the offending user sessions (2 session from APEX)
- 9:50 - cmsr3 up -> full cluster capacity restored
- 9:50 - conditions replication unstuck and fully functional -> streams processes restored full functionality automatically
- 11:20 - pvss replication unstuck and fully functional -> this required manual intervention
- 12:22 - cmsr3 reboots -> node rebooted by cluster, services relocated to remaining nodes.
- 12:30 cmsr4 reboots -> node rebooted by cluster, services relocated to remaining nodes.
- 12:50 - pvss replication unstuck and fully functional -> this required manual intervention
- 13:10 - condition replication unstuck and fully functional -> this required manual intervention

Analysis

The root cause of reboots is still not clear. OS watcher didn't help to diagnose it. Also the reason Streams replication was so badly affected is not clear.

Follow up

Since there is suspicion that node reboots are caused by excessive PGA utilization leading to swapping, an extra home-made PGA monitoring tool has been deployed on the cluster. Extra OCR tracking will be enabled, too.

-- EvaDafonte - 07-Jun-2010

This topic: DB > PostMortem20Aug10
Topic revision: r1 - 2010-08-25 - unknown



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
or Ideas, requests, problems regarding TWiki? use [Discourse](#) or [Send feedback](#)