

## Description

On Wednesday afternoon, at 2pm, the Atlas offline production database (ATLR) was affected by a scheduled intervention aimed at replacing a defective redundant FC switch. The intervention which was expected to be transparent (and in the past was performed few times in a transparent way) this time caused a sequence of serious issues and ended up with database unavailability and restart of all Oracle instances.

## Impact

- 2:00PM - 3:00PM: ATLR database was running with reduced number of nodes
- 3:00PM - 4:15PM: ATLR database was unavailable for the users' community

## Time line of the incident

All events described below happened on 15.12.2010.

- 2:00PM - Intervention started. FC switch was powered off and 2 nodes out of 5 of ATLR DB rebooted.
- 3:00PM - During the new FC switch reconfiguration some IO errors were observed and remaining instances went to hanging state.
- 3:45PM - Because ALTR database was unavailable and all instances were not responding, decision to restart them one by one were made.
- 4:15PM - ATLR database were back to operation and after some sanity checks database was available to the users' community

## Analysis

Reboots of 2 nodes were most probably caused by problems with accessing the cluster voting disks within configured timeout. However the main problem which showed up around 3PM is still the subject of investigation.

## Follow up

We've updated the procedures for future interventions of this type to be done in a set of steps of smaller scope to better control the risks. The hang of Oracle instances is being investigated with Oracle Support Services.

---

This topic: DB > PostMortem16Dec10

Topic revision: r2 - 2010-12-17 - MarcinBłaszczuk



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback