

GDP - Feature #48

Scanning for corrupt indices at gdplogd startup

08/23/2016 02:12 PM - Nitesh Mor

Status:	Closed	Start date:	08/23/2016
Priority:	Low	Due date:	
Assignee:	Eric Allman	% Done:	0%
Category:	gdplogd	Estimated time:	0.00 hour
Target version:			
Description			
<p>Even with appropriate flushing of appends to data file and to index files (recno=>offset and timestamp=>recno), there is a chance of application/system crash at a specific moment which will end up in out-of-sync condition for the index files.</p> <p>The root cause of this is that the update to data file and the index files is not an atomic operation. It'd be nice to have the ACID semantics for such updates, but one could argue that the index files could always be rebuilt from the data file, and hence the strong ACID semantics aren't needed for the entire operation. However, this leads to the necessity for detecting such corruption and repairing it.</p> <p>There are various ways to do it, here is just one possible way: The corruption is limited to only one record (assuming the data is flushed with each append individually). At startup time, gdplogd has to scan the local disk for the available records anyways. An extra step would be to verify the integrity of each index file by just looking at the most recent entry, and if needed, do a repair. Another optimization could be to keep a soft-state on persistent storage to mark non-corrupt logs on a clean shutdown.</p>			
Related issues:			
Related to GDP - Feature #46: Need way of rebuilding log indices ("fsck for l...		Closed	08/22/2016

History

#1 - 08/30/2016 11:11 AM - Nitesh Mor

- Related to Feature #46: Need way of rebuilding log indices ("fsck for logs") added

#2 - 04/09/2019 02:09 PM - Eric Allman

- Status changed from New to Closed

Not relevant for GDPv2, which uses SQLite for log storage.