

Ceph - Feature #3805

log: detect dup messages

01/15/2013 02:17 PM - Sage Weil

Status: New	% Done: 0%
Priority: Normal	Spent time: 0.00 hour
Assignee:	
Category: common	
Target version:	
Source:	Reviewed:
Tags:	Affected Versions:
Backport:	Pull request ID:
Description	
If a log message comes through and is a dup of the previous, increment a counter or something and only log it once with a (repeated N times) type message.	
Related issues:	
Related to Ceph - Feature #3775: log: stop logging in stats reports usage ab...	New

History

#1 - 01/15/2013 02:25 PM - Greg Farnum

What kind of dups are we trying to detect?

This sounds to me like a wishlist item that requires much more work to be useful than we'd like. From Deb's previous ticket comments I think she'd like to see similar outputs (but with different entities associated) being compressed down. But with most of those kinds of messages, the entity is the important part in the rare occasion when we care about the log message.

#2 - 01/15/2013 06:21 PM - Dan Mick

I tend to think there aren't very many dups we could usefully compress. It's pretty easy to add a one-string buffer to compare everything but the timestamp, but I suspect it also wouldn't compress very much. I left 3775 as "need more info" to suggest some strings that could profitably be removed or compressed because they have little useful info in them, but I suspect there aren't many.

#3 - 01/16/2013 07:03 AM - Sam Lang

The one that comes to mind is "no heartbeat from osd.foo since timestamp bar" messages. We could try to identify the few cases where this does happen (grep the mailing list maybe?), and add appropriate backoff/escalation logic to those cases. I suspect that repeat messages either can be ignored and should be summarized (as Sage suggests), such as clock skew, or are an indication of something severe (probably more severe as the number of messages increase), so handling them all in the same way might not be appropriate. In the severe cases, can we start reporting outside of logging (output in the ceph status summary, start sending messages to all the terminals on the node, etc.)?