

Ceph - Feature #19657

An EIO from a single device should not be a client-visible failure.

04/18/2017 06:58 PM - David Zafman

Status: Resolved	% Done: 0%
Priority: Urgent	Spent time: 0.00 hour
Assignee: David Zafman	
Category:	
Target version: v12.1.0	
Source:	Reviewed:
Tags:	Affected Versions:
Backport:	Pull request ID:
Description	
<p>Sage: the underlying goal: any EIO from a single device should not be a client-visible failure. (the same way an OSD failure shouldn't be visible to client workloads)</p> <p>we get this by enabling the eio on read config option, but the goal would be to do better than that: trigger a repair on an object that gives EIO on read, just like we do during scrub. (well, during repair)</p> <p>or, if we don't repair, work around it by reading from a replica that doesn't return eio. (which might make more sense since repair after scrub pbly won't be driven directly by the osd in the future?)</p> <p>davidz ok. yes... eventually repair is moving out of osd and we won't be using recovery to do it (hopefully). We could use that mechanism for now I guess. Or as you say may we just read from another replica</p>	
Related issues:	
Duplicated by RADOS - Bug #20863: CRC error does not mark PG as inconsistent ... Duplicate 07/31/2017	

History

#1 - 04/18/2017 08:20 PM - David Zafman

- Priority changed from Normal to Urgent

#2 - 04/18/2017 08:23 PM - David Zafman

- Target version set to v12.1.0

#3 - 05/23/2017 08:25 PM - David Zafman

- Status changed from New to In Progress

#4 - 06/29/2017 09:18 PM - David Zafman

<https://github.com/ceph/ceph/pull/14760>

#5 - 06/29/2017 09:19 PM - David Zafman

- Status changed from In Progress to Resolved

#6 - 08/08/2017 06:36 PM - David Zafman

- Duplicated by Bug #20863: CRC error does not mark PG as inconsistent or queue for repair added