

RADOS - Feature #44107

mon: produce stable election results when netsplits and other errors happen

02/12/2020 11:53 PM - Greg Farnum

Status:	Resolved	% Done:	0%
Priority:	Urgent	Spent time:	0.00 hour
Assignee:	Greg Farnum		
Category:			
Target version:			
Source:		Affected Versions:	
Tags:		Component(RADOS):	Monitor
Backport:	nautilus	Pull request ID:	35906
Reviewed:			
Description Right now, in netsplits and similar error conditions the monitors do not produce a stable quorum: whichever monitors are excluded will prompt continuous elections by Proposing to whatever peers they can reach. To accomplish this, add heartbeating between the monitor daemons, use that to generate connection liveness and reliability scores, and use those scores as input to an election algorithm. https://github.com/ceph/ceph/pull/32336			
Related issues: Blocks RADOS - Feature #44108: mon: osd: handle 2-(main-)site stretch cluster... In Progress			

History

#1 - 02/12/2020 11:56 PM - Greg Farnum

- Status changed from New to Fix Under Review

#2 - 02/12/2020 11:59 PM - Greg Farnum

- Blocks Feature #44108: mon: osd: handle 2-(main-)site stretch clusters explicitly, so no admin intervention is needed when a DC dies added

#3 - 02/13/2020 03:50 PM - Ken Dreyer

- Backport set to nautilus

#4 - 02/26/2020 09:50 PM - Neha Ojha

- Priority changed from Normal to Urgent

Marking anything we need for octopus as "Urgent".

#5 - 02/21/2022 05:09 PM - Greg Farnum

- Status changed from Fix Under Review to Resolved

- Pull request ID changed from 32336 to 35906

Oh, this has been done for ages.