

Ceph - Feature #9598

re-enable Objecter fast dispatch

09/25/2014 02:14 PM - Greg Farnum

Status:	Resolved	Start date:	09/25/2014
Priority:	Normal	Due date:	
Assignee:	Sage Weil	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	v0.93 - Last Hammer Sprint	Spent time:	0.00 hour
Source:	Development	Reviewed:	
Tags:		Affected Versions:	
Backport:	giant	Pull request ID:	
Description			
We had to nix fast dispatch on the Objecter because it could deadlock in conjunction with mark_down() calls.			
Fixing this will require either: 1) changing the mark_down() and associated interface requirements so the messenger doesn't block (possible, I think? we'd need to be careful about how Dispatchers handle their notifications to avoid introducing new deadlocks) 2) reworking the Objecter locking around OSDSessions a little.			
Whoever does this should explore both options to see which is most useful and minimally invasive. See #9462 and #9295 for some history.			
Related issues:			
Related to Messengers - Bug #9898: osd: fast dispatch deadlock in mark_down (...)		Resolved	10/26/2014

Associated revisions

Revision c9f9e72e - 10/14/2014 07:42 PM - Sage Weil

Revert "Objecter: disable fast dispatch of CEPH_MSG_OSD_OPREPLY messages"

This reverts commit 3f23709c474292f9239f77a6cce26309fc86ce29.

We have changed mark_down() behavior so that it no longer blocks on fast dispatch. This makes the objecter reply handler safe again.

Fixes: #9598

Signed-off-by: Sage Weil <sage@redhat.com>

Revision 1261bf24 - 01/12/2015 07:17 PM - Sage Weil

Revert "Objecter: disable fast dispatch of CEPH_MSG_OSD_OPREPLY messages"

This reverts commit 3f23709c474292f9239f77a6cce26309fc86ce29.

We have changed mark_down() behavior so that it no longer blocks on fast dispatch. This makes the objecter reply handler safe again.

Fixes: #9598

Signed-off-by: Sage Weil <sage@redhat.com>

(cherry picked from commit c9f9e72e558521cb90f90538bc27f995f82d76c2)

History

#1 - 09/30/2014 01:21 PM - Ian Colle

- Assignee set to Sage Weil

#2 - 10/14/2014 12:43 PM - Sage Weil

- Target version changed from 0.88 to 0.89

#3 - 10/14/2014 01:24 PM - Sage Weil

- Status changed from New to Need Review

#4 - 10/27/2014 01:38 PM - Sage Weil

- Status changed from Need Review to Testing

- Target version changed from 0.89 to 0.88

#5 - 11/11/2014 09:36 AM - Sage Weil

- Target version changed from 0.88 to 0.90

#6 - 11/11/2014 01:07 PM - Samuel Just

- Target version changed from 0.90 to v.91

#7 - 11/12/2014 10:00 AM - Sage Weil

- Status changed from Testing to Pending Backport

sage-2014-11-11_08:26:01-rados-wip-sage-testing-distro-basic-multi

#8 - 11/25/2014 01:03 PM - Samuel Just

- Target version changed from v.91 to v.actually90

#9 - 11/25/2014 01:03 PM - Samuel Just

- Target version changed from v.actually90 to v.91

#10 - 12/09/2014 05:29 AM - Ian Colle

- Target version changed from v.91 to v.actually90

#11 - 01/12/2015 02:09 PM - Loic Dachary

- giant backport <https://github.com/ceph/ceph/pull/3356>

#12 - 01/12/2015 02:10 PM - Loic Dachary

- Backport set to giant

#13 - 01/20/2015 09:25 AM - Ian Colle

- Target version changed from v.actually90 to v0.93 - Last Hammer Sprint

#14 - 02/10/2015 09:57 PM - Loic Dachary

- Status changed from Pending Backport to Resolved