

Ceph - Bug #10257

Ceph df doesn't report MAX AVAIL correctly when using rulesets and OSD in ruleset is down and out

12/05/2014 12:08 PM - Xavier Trilla

Status:	Resolved	% Done:	0%
Priority:	High	Spent time:	0.00 hour
Assignee:	Joao Eduardo Luis		
Category:	Monitor		
Target version:			
Source:	Community (user)	Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:	giant,firefly	Pull request ID:	
Regression:	No	Crash signature (v1):	
Severity:	3 - minor	Crash signature (v2):	
Reviewed:			

Description

In our setup we have two rulesets, one for SSDs and another one for HDDs. Ceph df normally reports the MAX AVAIL space considering the OSDs in the ruleset, but when one of the OSDs is down and out it just reports 0 instead of the real MAX AVAIL space for the pools using that ruleset.

Command "ceph fs" output:

GLOBAL:

SIZE	AVAIL	RAW USED	%RAW USED	OBJECTS
28309G	22750G	5558G	19.63	757k

POOLS:

NAME	ID	CATEGORY	USED	%USED	MAX AVAIL	OBJECTS	DIRTY	READ	WRITE
ssd-volumes-01	118	-	5344M	0.02	0	1472	1472	4861k	5281k
hdd-volumes-01	120	-	0	0	8200G	0	0	0	0

ssd-volumes-01 pool is using ruleset 0 (It has one OSD down and out)

hdd-volumes-01 pool is using ruleset 1 (All OSDs are up)

Related issues:

Related to Ceph - Bug #13840: Ceph Pools' MAX AVAIL is 0 if some OSDs' weight...

Resolved

11/20/2015

Associated revisions

Revision b311e7c3 - 01/20/2015 04:18 PM - Joao Eduardo Luis

mon: PGMonitor: skip zeroed osd stats on get_rule_avail()

Fixes: #10257

Signed-off-by: Joao Eduardo Luis <joao@redhat.com>

Revision 6e58732f - 02/27/2015 05:39 PM - Joao Eduardo Luis

mon: PGMonitor: skip zeroed osd stats on get_rule_avail()

Fixes: #10257

Signed-off-by: Joao Eduardo Luis <joao@redhat.com>

(cherry picked from commit b311e7c36273efae39aa2602c1f8bd90d39e5975)

Conflicts:

src/mon/PGMonitor.cc

ceph::unordered_map changed the context, simple resolution

Revision a5a76f81 - 03/17/2015 09:55 AM - Joao Eduardo Luis

mon: PGMonitor: skip zeroed osd stats on get_rule_avail()

Fixes: #10257

Signed-off-by: Joao Eduardo Luis <joao@redhat.com>
(cherry picked from commit b311e7c36273efae39aa2602c1f8bd90d39e5975)

History

#1 - 12/08/2014 05:12 AM - Zheng Yan

- Project changed from CephFS to Ceph
- Target version deleted (v0.87)

#2 - 12/08/2014 09:50 AM - Sage Weil

- Category set to Monitor
- Priority changed from Normal to High
- Source changed from other to Community (user)
- Tags deleted (ruleset, df, down, out)

#3 - 12/16/2014 01:26 PM - Sage Weil

- Assignee set to Joao Eduardo Luis

#4 - 01/20/2015 11:29 AM - Sage Weil

- Status changed from New to Pending Backport
- Backport set to giant,firefly

#5 - 01/20/2015 11:39 AM - Joao Eduardo Luis

to be backported to giant and firefly: 8be6a6ab2aa5a000a39c73a98b11a0ab32ffa1c and b311e7c36273efae39aa2602c1f8bd90d39e5975.
50547dc3c00b7556e26b9a44ec68640c5c3a2384 needs to be backported to dumpling too.

#6 - 02/27/2015 05:25 PM - Loïc Dachary

- Backport changed from giant,firefly to giant,firefly,dumpling

#7 - 03/17/2015 09:59 AM - Loïc Dachary

- giant backport <https://github.com/ceph/ceph/pull/4027>

#8 - 03/17/2015 05:42 PM - Sage Weil

giant and firefly backports done, still missing dumpling

#9 - 03/26/2015 02:40 PM - Loïc Dachary

- Backport changed from giant,firefly,dumpling to giant,firefly

dumpling is end of life

#10 - 03/26/2015 02:41 PM - Loïc Dachary

[a5a76f8](#) mon: PGMonitor: skip zeroed osd stats on get_rule_avail() (in giant), [6e58732](#) mon: PGMonitor: skip zeroed osd stats on get_rule_avail() (in firefly),

#11 - 03/26/2015 02:41 PM - Loïc Dachary

- Status changed from Pending Backport to Resolved

#12 - 11/20/2015 10:21 AM - Nathan Cutler

- Related to Bug #13840: Ceph Pools' MAX AVAIL is 0 if some OSDs' weight is 0 added