

rbd - Bug #38538

Performance improvements for object-map

03/01/2019 03:27 PM - Jason Dillaman

Status:	Pending Backport	% Done:	0%
Priority:	Normal	Spent time:	0.00 hour
Assignee:	Jason Dillaman		
Category:			
Target version:			
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:	luminous,mimic	ceph-qa-suite:	
Regression:	No	Pull request ID:	26721
Severity:	3 - minor	Crash signature:	
Description			
<p>The "object_map_update" OSD class method has a some low-hanging improvements that can be made to increase the performance of object-map updates:</p> <p>1) the std::vector BitVector used for the bitset data is zeroed before it's contents are replaced with the bitset data read from the object-map object. This is a huge unnecessary CPU hit for larger object-maps.</p> <p>2) the entire footer is read into memory even though only 4 bytes might be changed for an updated CRC. For large RBD images, this could be up to 16KiB of data in the worst-case scenario.</p> <p>A stress-test against librados_test_stub shows improvement from ~100 updates/sec to over 50K updates/sec. The OSD is still limited to approximately 2-3K write operations per second, so longer term refactoring of the object-map design will be required to improve it's performance by orders of magnitude.</p>			
Related issues:			
Copied to rbd - Backport #38673: mimic: Performance improvements for object-map		In Progress	
Copied to rbd - Backport #38674: luminous: Performance improvements for objec...		Resolved	

History

#1 - 03/01/2019 04:23 PM - Jason Dillaman

- Status changed from In Progress to Fix Under Review
- Pull request ID set to 26721

#2 - 03/11/2019 02:45 PM - Mykola Golub

- Status changed from Fix Under Review to Pending Backport

#3 - 03/11/2019 02:54 PM - Nathan Cutler

- Copied to Backport #38673: mimic: Performance improvements for object-map added

#4 - 03/11/2019 02:54 PM - Nathan Cutler

- Copied to Backport #38674: luminous: Performance improvements for object-map added

#5 - 09/17/2019 11:14 PM - Jason Dillaman

- Description updated