

CephFS - Feature #58550

Feature # 58488 (New): mds: avoid encoding snode for each ancestor in an EMetaBlob log event

mds: add perf counter to track (relatively) larger log events

01/23/2023 03:12 PM - Venky Shankar

Status:	Resolved	% Done:	0%
Priority:	Normal		
Assignee:	Patrick Donnelly		
Category:	Administration/Usability		
Target version:	v19.0.0		
Source:	Development	Affected Versions:	
Tags:		Component(FS):	MDS
Backport:		Labels (FS):	
Reviewed:		Pull request ID:	48732
Description			
Logging of large (in size) subtreemap over and over again in the log segment is a well know scalability limiting factor in the MDS. Its a bit hard to infer this though when a cluster starts experiencing slowness. Moreover, there <i>might</i> be other places where a large log event size could be making things worse. Adding a perf counter that increments whenever the log event size exceeds a certain threshold might help is inferring the extent of logging large log events over and over again.			

History

#1 - 01/30/2023 01:53 PM - Venky Shankar

- Category set to Administration/Usability
- Assignee set to Patrick Donnelly

#2 - 02/01/2023 08:14 PM - Patrick Donnelly

- Status changed from New to In Progress

#3 - 02/01/2023 11:34 PM - Patrick Donnelly

- Status changed from In Progress to Fix Under Review
- Pull request ID set to 48732

#4 - 08/07/2023 03:34 PM - Patrick Donnelly

- Status changed from Fix Under Review to Resolved
- Target version changed from v18.0.0 to v19.0.0
- Source set to Development
- Backport deleted (pacific,quincy)

Not planning to backport this the minor log segment PR.