

fs - Feature #25013

mds: add average session age (uptime) perf counter

07/20/2018 06:14 AM - Venky Shankar

Status:	Resolved	Start date:	07/20/2018
Priority:	Normal	Due date:	
Assignee:	Venky Shankar	% Done:	0%
Category:	Administration/Usability	Estimated time:	0.00 hour
Target version:	v14.0.0	Affected Versions:	
Source:	Development	Component(FS):	MDS
Tags:		Labels (FS):	task(medium)
Backport:	mimic,luminous	Pull request ID:	
Reviewed:			

Description

Add a performance counter in the MDS to capture average age (uptime) of sessions. This should reflect the average uptime of ongoing and closed sessions (as a single perf stat). This was marked to be included in PR <https://github.com/ceph/ceph/pull/22668> initially, but had to be carved out as a separate feature ticket due to pending discussions about the best method (in terms of correctness and efficiency) to calculate this stat.

Thoughts from earlier discussions:

You can add a public method to get the average session starttime: keep a running average of the starttimes (the co then divide by the number of sessions when computing the average. When adding a session:

$$\text{avg_starttime} = ((\text{avg_starttime} / (\text{new_session_count})) * (\text{new_session_count}-1) + \text{sesion_starttime}/\text{new_session_count}$$

Removing is the reverse.

The tricky part is doing that with monotonic clock time points rather than unix epochs. @adamemerson, can you advise the right way to compute the average of a set of coarse_mono_time epochs? We need to be able to add/remove epochs from the set dynamically, using the math above ^

@batrick -- I think the average session age would need to be updated regularly (on each tick()) and not only on session add and remove. That way we would have the average uptime for ongoing sessions too which would be constantly refreshed (but is computationally expensive if there are huge number of sessions).

Related issues:

Related to fs - Bug #24004: mds: curate priority of perf counters sent to mgr	Resolved	05/03/2018
Related to fs - Bug #35411: mds: store session birth time on-disk in session map	In Progress	09/03/2018
Copied to fs - Backport #35937: luminous: mds: add average session age (uptim...	Resolved	
Copied to fs - Backport #35938: mimic: mds: add average session age (uptime) ...	Resolved	

History

#1 - 07/20/2018 06:15 AM - Venky Shankar

- Target version set to v14.0.0

- Source set to Development

#2 - 07/20/2018 06:16 AM - Venky Shankar

- Related to Bug #24004: mds: curate priority of perf counters sent to mgr added

#3 - 07/20/2018 10:23 PM - Patrick Donnelly

- Status changed from New to In Progress

- Assignee set to Venky Shankar

#4 - 07/30/2018 09:25 AM - Venky Shankar

- Status changed from In Progress to Need Review

PR <https://github.com/ceph/ceph/pull/23314>

#5 - 09/03/2018 04:29 AM - Venky Shankar

- Related to Bug #35411: mds: store session birth time on-disk in session map added

#6 - 09/11/2018 06:22 PM - Patrick Donnelly

- Status changed from Need Review to Pending Backport

#7 - 09/11/2018 07:16 PM - Nathan Cutler

- Copied to Backport #35937: luminous: mds: add average session age (uptime) perf counter added

#8 - 09/11/2018 07:16 PM - Nathan Cutler

- Copied to Backport #35938: mimic: mds: add average session age (uptime) perf counter added

#9 - 11/13/2018 09:57 AM - Nathan Cutler

- Status changed from Pending Backport to Resolved