

mgr - Feature #38211

mgr/dashboard: Add custom dialogue for configuring PG scrub parameters

02/06/2019 12:47 PM - Lenz Grimmer

Status: Resolved	% Done: 0%
Priority: Normal	
Assignee: Tatjana Dehler	
Category: dashboard/osds	
Target version: v15.0.0	
Source:	Reviewed:
Tags: dashboard, management	Affected Versions: v14.0.0, v14.2.0, v14.2.1, v15.0.0
Backport: nautilus	Pull request ID: 27072
Description	
<p>Ceph provides multiple configuration options to tweak the behavior of scrub operations, e.g.</p> <ul style="list-style-type: none">• <code>osd_max_scrubs</code> - the maximum number of simultaneous scrub operations on a given OSD• <code>osd_scrub_during_recovery</code> - Allow scrubbing when PGs on the OSD are undergoing recovery• <code>osd_scrub_begin_hour</code>, <code>osd_scrub_end_hour</code> - the hours of day (0 to 24) that define a time window when the scrubbing can happen• <code>osd_scrub_begin_week_day</code>, <code>osd_scrub_end_week_day</code> - the days of week (0 to 6) that define a time window when the scrubbing can happen• <code>osd_scrub_min_interval</code>, <code>osd_scrub_max_interval</code> - Scrub each PG no more/no less often than this interval• <code>osd_deep_scrub_interval</code> - Deep scrub each PG (i.e., verify data checksums) at least this often• <code>osd_scrub_auto_repair_num_errors</code> - Maximum number of detected errors to automatically repair <p>While it's possible to modify these via the config editor, it may be worthwhile providing a custom "Scrub configuration" dialogue, similar to the recovery profiles or the global OSD settings. Are there any other settings that should be added here?</p>	
Related issues:	
Copied to mgr - Backport #40059: nautilus: mgr/dashboard: Add custom dialogue... Resolved	

History

#1 - 02/18/2019 03:38 PM - Tatjana Dehler

Lenz Grimmer wrote:

Are there any other settings that should be added here?

By 'settings' you mean any other config options, don't you?

I had a look at the config options page and found 32 OSD options where 'scrub' is mentioned:

- `osd_debug_deep_scrub_sleep` -> Inject an expensive sleep during deep scrub IO to make it easier to induce preemption
- `osd_deep_scrub_interval` -> Deep scrub each PG (i.e., verify data checksums) at least this often
- `osd_deep_scrub_keys` -> Number of keys to read from an object at a time during deep scrub
- `osd_deep_scrub_large_omap_object_key_threshold` -> Warn when we encounter an object with more omap keys than this
- `osd_deep_scrub_large_omap_object_value_sum_threshold` -> Warn when we encounter an object with more omap key bytes than this
- `osd_deep_scrub_randomize_ratio` -> Scrubs will randomly become deep scrubs at this rate (0.15 -> 15% of scrubs are deep)
- `osd_deep_scrub_stride` -> Number of bytes to read from an object at a time during deep scrub
- `osd_deep_scrub_update_digest_min_age` -> Update overall object digest only if object was last modified longer ago than this
- `osd_max_scrubs` -> Maximum concurrent scrubs on a single OSD
- `osd_op_queue_mclock_scrub_lim` -> mclock weight of limit requests
- `osd_op_queue_mclock_scrub_res` -> mclock reservation of scrub requests
- `osd_op_queue_mclock_scrub_wgt` -> mclock weight of scrub requests
- `osd_requested_scrub_priority`
- `osd_scrub_auto_repair` -> Automatically repair damaged objects detected during scrub
- `osd_scrub_auto_repair_num_errors` -> Maximum number of detected errors to automatically repair
- `osd_scrub_backoff_ratio` -> Backoff ratio for scheduling scrubs
- `osd_scrub_begin_hour` -> Restrict scrubbing to this hour of the day or later

- **osd_scrub_begin_week_day** -> Restrict scrubbing to this day of the week or later
- **osd_scrub_chunk_max** -> Maximum number of objects to scrub in a single chunk
- **osd_scrub_chunk_min** -> Minimum number of objects to scrub in a single chunk
- **osd_scrub_cost** -> Cost for scrub operations in work queue
- **osd_scrub_during_recovery** -> Allow scrubbing when PGs on the OSD are undergoing recovery
- **osd_scrub_end_hour** -> Restrict scrubbing to hours of the day earlier than this
- **osd_scrub_end_week_day** -> Restrict scrubbing to days of the week earlier than this
- **osd_scrub_interval_randomize_ratio** -> Ratio of scrub interval to randomly vary
- **osd_scrub_invalid_stats**
- **osd_scrub_load_threshold** -> Allow scrubbing when system load divided by number of CPUs is below this value
- **osd_scrub_max_interval** -> Scrub each PG no less often than this interval
- **osd_scrub_max_preemptions** -> Set the maximum number of times we will preempt a deep scrub due to a client operation before blocking IO to complete the scrub
- **osd_scrub_min_interval** -> Scrub each PG no more often than this interval
- **osd_scrub_priority** -> Priority for scrub operations in work queue
- **osd_scrub_sleep** -> Duration to inject a delay during scrubbing

#2 - 02/19/2019 08:12 AM - Tatjana Dehler

- Assignee set to *Tatjana Dehler*

#3 - 02/19/2019 11:50 AM - Lenz Grimmer

- Status changed from *New* to *In Progress*

#4 - 03/31/2019 11:31 PM - Lenz Grimmer

- Target version set to *v15.0.0*

- Backport set to *nautilus*

#5 - 04/04/2019 06:47 AM - Tatjana Dehler

- Pull request ID set to *27072*

#6 - 05/13/2019 01:47 PM - Patrick Seidensal

- Status changed from *In Progress* to *Fix Under Review*

#7 - 05/29/2019 08:25 AM - Lenz Grimmer

- Status changed from *Fix Under Review* to *Pending Backport*

- Affected Versions *v14.0.0, v14.2.0, v14.2.1, v15.0.0* added

#8 - 05/29/2019 11:40 AM - Nathan Cutler

- Copied to Backport #40059: *nautilus: mgr/dashboard: Add custom dialogue for configuring PG scrub parameters* added

#9 - 06/21/2019 08:06 AM - Lenz Grimmer

- Status changed from *Pending Backport* to *Resolved*