

Orchestrator - Bug #47873

/usr/lib/sysctl.d/90-ceph-osd.conf getting installed in container, rendering it ineffective

10/15/2020 06:40 PM - Nathan Cutler

| | |
|---|-------------------------------|
| Status: Resolved | % Done: 0% |
| Priority: High | |
| Assignee: Michael Fritch | |
| Category: cephadm (binary) | |
| Target version: | |
| Source: | Affected Versions: |
| Tags: low-hanging-fruit | ceph-qa-suite: |
| Backport: | Pull request ID: 41264 |
| Regression: No | Crash signature (v1): |
| Severity: 3 - minor | Crash signature (v2): |
| Reviewed: | |
| Description | |
| The file /usr/lib/sysctl.d/90-ceph-osd.conf has the following contents: | |
| <pre>fs.aio-max-nr = 1048576 kernel.pid_max = 4194304</pre> | |
| This is very important to have on all OSD nodes. However, since this file is owned by the ceph-osd package, and that package is installed in the container image, and not in the host system, it no longer takes effect. This causes problems on OSDs nodes with larger numbers of disks. | |
| One option would be for cephadm to manage this file, ensuring that it gets created on OSD nodes when OSDs are created, and removing it when a node no longer has any OSDs. | |
| Another option would be to move the file to ceph-base. Then it would be installed (in the host system) on all cluster nodes, regardless of whether the node has OSDs or not. | |
| Yet another option would be to move the file into ceph-common. In this scenario, the file would be installed on both cluster nodes and client nodes as well. | |
| Related issues: | |
| Related to Orchestrator - Feature #49171: cephadm: set osd-memory-target | In Progress |
| Related to Orchestrator - Bug #50551: Massive OSD creation: kernel parameter ... | Duplicate |
| Related to Orchestrator - Feature #48292: cephadm: allow more than 60 OSDs pe... | New |
| Related to Orchestrator - Bug #50526: OSD massive creation: OSDs not created | Resolved |

History

#1 - 10/16/2020 06:01 AM - Nathan Cutler

- Subject changed from /usr/lib/sysctl.d/90-ceph-osd.conf getting installed in container to /usr/lib/sysctl.d/90-ceph-osd.conf getting installed in container, rendering it ineffective

#2 - 11/19/2020 10:47 AM - Sebastian Wagner

- Priority changed from Normal to High

#3 - 02/04/2021 04:01 PM - Sebastian Wagner

- Related to Feature #49171: cephadm: set osd-memory-target added

#4 - 03/16/2021 08:58 AM - Sebastian Wagner

- Tags set to *low-hanging-fruit*

#5 - 04/28/2021 09:26 AM - Sebastian Wagner

- Related to Bug #50551: Massive OSD creation: kernel parameter *fs.aio-max-nr* with a low value by default added

#6 - 04/28/2021 09:28 AM - Sebastian Wagner

- Related to Feature #48292: *cephadm*: allow more than 60 OSDs per host added

#7 - 04/29/2021 09:01 AM - Juan Miguel Olmo Martínez

- Related to Bug #50526: OSD massive creation: OSDs not created added

#8 - 05/10/2021 08:09 PM - Michael Fritch

- Category set to *cephadm* (binary)

- Assignee set to Michael Fritch

- Pull request ID set to 41264

#9 - 05/20/2021 06:25 AM - Kefu Chai

- Status changed from *New* to *Resolved*