

Linux kernel client - Feature #43215

kclient: send usage metrics to MDS

12/09/2019 06:52 PM - Patrick Donnelly

Status: Resolved	% Done: 100%
Priority: Normal	Spent time: 0.00 hour
Assignee: Xiubo Li	
Category:	
Target version: v16.0.0	
Source: Development	Reviewed:
Tags:	Affected Versions:
Backport:	
Description kernel client version of #36253	
Subtasks:	
Feature # 44533: kclient: add min/max perf latency support	Resolved
Feature # 44534: kclient: add standard deviation support for the read/write/metadata la...	Resolved
Related issues:	
Related to CephFS - Feature #36253: cephfs: clients should send usage metadat...	Resolved
Related to CephFS - Feature #24285: mgr: add module which displays current us...	Resolved
Related to CephFS - Feature #43423: mds: collect and show the dentry lease me...	Resolved
Related to CephFS - Feature #43435: kclient:send client provided metric flags...	Resolved

History

#1 - 12/09/2019 06:52 PM - Patrick Donnelly

- Related to Feature #36253: cephfs: clients should send usage metadata to MDSs for administration/monitoring added

#2 - 12/09/2019 06:52 PM - Patrick Donnelly

- Related to Feature #24285: mgr: add module which displays current usage of file system (`fs top`) added

#3 - 12/14/2019 12:52 AM - Xiubo Li

- Status changed from New to In Progress

#4 - 12/16/2019 12:01 PM - Xiubo Li

For both dentry lease and caps perf metric we will only count the hit/miss info triggered from the vfs calls, for the cases like request reply handling and periodically ceph_trim_dentries() we will ignore them.

Currently only the debugfs is support and next will fulfill sending the mertic data to MDS.

The output will be:

```
# cat /sys/kernel/debug/ceph/536b8f9c-0d78-46bc-bbe9-c2ceb61c82a0.client4275/metric
item          hit          miss
-----
d_lease       19           0
i_caps        168          1
```

#5 - 12/17/2019 12:21 PM - Xiubo Li

Will add the entry count and cap count support:

The output will be:

item	total	hit	miss
d_lease	11	141	7
i_caps	10	28995740	13

#6 - 12/26/2019 06:16 AM - Xiubo Li

- Related to Feature #43423: mds: collect and show the dentry lease metric added

#7 - 12/26/2019 06:21 AM - Xiubo Li

The output will be:

```
# cat /sys/kernel/debug/ceph/02a220c4-e4c6-43cd-9a84-993a020d2fffb.client4273/metrics
item          total          miss           hit
-----
d_lease       125            120            2184

session       caps           miss           hit
-----
0             63            692            7100
1             58            0              113
2             66            0              109
```

```
# ./bin/ceph fs perf stats | python -m json.tool
{
  "client_metadata": {
    "client.4273": {
      "IP": "v1:192.168.195.165",
      "hostname": "federal",
      "mount_point": "N/A",
      "root": "/"
    }
  },
  "counters": [
    "cap_hit"
  ],
  "global_counters": [
    "read_latency",
    "write_latency",
    "metadata_latency",
    "dentry_lease_hit"
  ],
  "global_metrics": {
    "client.4273": [
      [
        0,
        0
      ],
      [
        0,
        0
      ],
      [
        0,
        0
      ],
      [
        2184,

```

```
120
  ]
]
},
"metrics": {
  "delayed_ranks": [],
  "mds.0": {
    "client.4273": [
      [
        7100,
        692
      ]
    ]
  },
  "mds.1": {
    "client.4273": [
      [
        113,
        0
      ]
    ]
  },
  "mds.2": {
    "client.4273": [
      [
        109,
        0
      ]
    ]
  }
}
}
```

#8 - 12/30/2019 01:57 AM - Xiubo Li

- Status changed from In Progress to Fix Under Review

#9 - 12/30/2019 05:44 AM - Xiubo Li

- Related to Feature #43435: kclient:send client provided metric flags in client metadata added

#10 - 12/30/2019 05:44 AM - Xiubo Li

- Related to Bug #24823: mds: deadlock when setting config value via admin socket added

#11 - 12/30/2019 05:47 AM - Xiubo Li

- Related to deleted (Bug #24823: mds: deadlock when setting config value via admin socket)

#12 - 12/30/2019 05:47 AM - Xiubo Li

- Related to Feature #43435: kclient:send client provided metric flags in client metadata added

#13 - 12/30/2019 05:47 AM - Xiubo Li

- Related to deleted (Feature #43435: kclient:send client provided metric flags in client metadata)

#14 - 01/06/2020 09:51 AM - Xiubo Li

Add read/write latency support:

```
# ./bin/ceph fs perf stats | python -m json.tool
{
  "client_metadata": {
    "client.4836": {
      "IP": "v1:192.168.195.165",
      "hostname": "federal",
      "mount_point": "N/A",
      "root": "/"
    }
  },
  "counters": [
    "cap_hit"
  ],
  "global_counters": [
    "read_latency",
    "write_latency",
    "metadata_latency",
    "dentry_lease_hit"
  ],
  "global_metrics": {
    "client.4836": [
      [
        0,
        15000000
      ],
      [
        13,
        77000000
      ],
      [
        0,
        0
      ],
      [
        6,
        0
      ]
    ]
  },
  "metrics": {
    "delayed_ranks": [],
    "mds.0": {
      "client.4836": [
        [
          1222,
          9
        ]
      ]
    },
    "mds.1": {
      "client.4836": [
```

```
    [
      4,
      0
    ]
  ]
}
}
```

#15 - 02/18/2020 09:55 PM - Patrick Donnelly

- Target version changed from v15.0.0 to v16.0.0

#16 - 03/10/2020 12:51 AM - Xiubo Li

For the patch series [1] it only support the debugfs showing, will add sending the metric to ceph stuff after the depending ceph PR [2] get merged.

[1]: <https://patchwork.kernel.org/project/ceph-devel/list/?series=253159>

[2]: <https://github.com/ceph/ceph/pull/26004>

Thanks

#17 - 06/18/2020 09:18 AM - Xiubo Li

V1: <https://patchwork.kernel.org/project/ceph-devel/list/?series=303647>

#18 - 07/23/2020 10:45 AM - Xiubo Li

- Project changed from CephFS to Linux kernel client

- Category deleted (Performance/Resource Usage)

- Status changed from Fix Under Review to Resolved

<https://patchwork.kernel.org/project/ceph-devel/list/?series=319465>

Merged, closing it.