

CephFS - Bug #44680

mds/Mutation.h: 128: FAILED ceph_assert(num_auth_pins == 0)

03/19/2020 07:12 PM - Sage Weil

Status: Resolved	% Done: 0%
Priority: Urgent	
Assignee: Zheng Yan	
Category:	
Target version:	
Source:	Affected Versions:
Tags:	ceph-qa-suite:
Backport: mimic, nautilus, octopus	Component(FS):
Regression: No	Labels (FS):
Severity: 3 - minor	Pull request ID: 34110
Reviewed:	Crash signature:

Description

```
{
  "assert_condition": "num_auth_pins == 0",
  "assert_file": "/home/jenkins-build/build/workspace/ceph-dev-build/ARCH/x86_64/AVAILABLE_ARCH/x86_64/AVAILABLE_DIST/centos8/DIST/centos8/MACHINE_SIZE/gigantic/release/15.1.1-91-g126c444/rpm/el8/BUILD/ceph-15.1.1-91-g126c444/src/mds/Mutation.h",
  "assert_func": "virtual MutationImpl::~MutationImpl()",
  "assert_line": 128,
  "assert_msg": "/home/jenkins-build/build/workspace/ceph-dev-build/ARCH/x86_64/AVAILABLE_ARCH/x86_64/AVAILABLE_DIST/centos8/DIST/centos8/MACHINE_SIZE/gigantic/release/15.1.1-91-g126c444/rpm/el8/BUILD/ceph-15.1.1-91-g126c444/src/mds/Mutation.h: In function 'virtual MutationImpl::~MutationImpl()' thread 7f061e8a0700 time 2020-03-19T02:09:57.010882+0000\n/home/jenkins-build/build/workspace/ceph-dev-build/ARCH/x86_64/AVAILABLE_ARCH/x86_64/AVAILABLE_DIST/centos8/DIST/centos8/MACHINE_SIZE/gigantic/release/15.1.1-91-g126c444/rpm/el8/BUILD/ceph-15.1.1-91-g126c444/src/mds/Mutation.h: 128 : FAILED ceph_assert(num_auth_pins == 0)\n",
  "assert_thread_name": "MR_Finisher",
  "backtrace": [
    "((+0x12dc0) [0x7f062ba3adc0]",
    "(pthread_getname_np()+0x48) [0x7f062ba3c038]",
    "(ceph::logging::Log::dump_recent()+0x428) [0x7f062cf49b28]",
    "((+0x4ab4cb) [0x560a9267a4cb]",
    "((+0x12dc0) [0x7f062ba3adc0]",
    "(gsignal()+0x10f) [0x7f062a4fe8df]",
    "(abort()+0x127) [0x7f062a4e8cf5]",
    "(ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x1a9) [0x7f062cc05c21]",
    "((+0x27adea) [0x7f062cc05dea]",
    "(MutationImpl::~MutationImpl()+0x205) [0x560a923e0e55]",
    "(TrackedOp::put()+0x71) [0x560a923c9a91]",
    "(C_Locker_FileUpdate_finish::~C_Locker_FileUpdate_finish()+0x32) [0x560a924e58a2]",
    "(MDSIOContextBase::complete(int)+0xfa) [0x560a925e6c7a]",
    "(MDSLogContextBase::complete(int)+0x44) [0x560a925e7044]",
    "(Finisher::finisher_thread_entry()+0x1a5) [0x7f062cc96385]",
    "((+0x82de) [0x7f062ba302de]",
    "(clone()+0x43) [0x7f062a5c3133]"
  ],
  "ceph_version": "15.1.1-91-g126c444",
  "crash_id": "2020-03-19T02:09:57.042400Z_0d2ae68f-e51b-4eee-8baa-b6186684d079",
  "entity_name": "mds.cephfs.reesi002.euduff",
  "os_id": "centos",
  "os_name": "CentOS Linux",
  "os_version": "8 (Core)",
}
```

```
"os_version_id": "8",
"process_name": "ceph-mds",
"stack_sig": "87c07aac5002b1b764575dc3e6e6411c5eac461da93d0097c1a0f4ef3d1bfd5e",
"timestamp": "2020-03-19T02:09:57.042400Z",
"utsname_hostname": "reesi002",
"utsname_machine": "x86_64",
"utsname_release": "4.4.0-116-generic",
"utsname_sysname": "Linux",
"utsname_version": "#140-Ubuntu SMP Mon Feb 12 21:23:04 UTC 2018"
}
```

two instances of this on the lab cluster this morning upgrading from yesterday's octopus to today's.

Related issues:

Related to CephFS - Bug #44295: mds: MDCache.cc: 6400: FAILED ceph_assert(r =...	Resolved
Copied to CephFS - Backport #45026: mimic: mds/Mutation.h: 128: FAILED ceph_a...	Rejected
Copied to CephFS - Backport #45027: nautilus: mds/Mutation.h: 128: FAILED cep...	Resolved
Copied to CephFS - Backport #45028: octopus: mds/Mutation.h: 128: FAILED ceph...	Resolved

History

#1 - 03/19/2020 08:49 PM - Greg Farnum

Do we have any logs or more detail about what happened?

The only thing this flags in my head is <https://github.com/ceph/ceph/pull/33291>, but that's in the Migrator.

Or Patrick merged a commit changing how we handle Contexts on shutdown a little bit which I was hinky about, but no real solid evidence.

#2 - 03/19/2020 08:56 PM - Greg Farnum

[13:55:18] <@sage> it was triggered by the upgrade... i'm guessing when the old container was stopped and got blacklisted?

[13:55:55] <@sage> i almost didn't notice because every upgrade i've been seeing 2 crashes on the lab cluster due to the blacklist error code from rados triggering an assert. but iirc that is fixed/cleaned up now

Okay on shutdown all I can think of is Context shutdown handling then.

#3 - 03/19/2020 09:02 PM - Greg Farnum

Yeah definitely the fault of <https://github.com/ceph/ceph/pull/33538>, which was trying to prevent us from asserting on EBLACKLIST errors on shutdown. But simply deletes any pending MDSIOContextBase on shutdown instead of letting them complete.

#4 - 03/19/2020 09:04 PM - Greg Farnum

- Assignee set to Zheng Yan

#5 - 03/23/2020 09:11 AM - Zheng Yan

maybe revert

<https://github.com/ceph/ceph/commit/73436961512bd87981244fa48212085faf7028c4> and
<https://github.com/ceph/ceph/commit/79f5052a1ddb61043de4e1cbe19ede2a6b4f53b>

and

don't assert empty io context list when shutting down

#6 - 03/23/2020 09:23 AM - Zheng Yan

- Status changed from New to Fix Under Review

- Pull request ID set to 34110

#7 - 03/24/2020 02:06 AM - Zheng Yan

- Backport set to octopus

#8 - 04/03/2020 07:38 PM - Greg Farnum

- Status changed from Fix Under Review to Pending Backport

#9 - 04/06/2020 08:24 PM - Nathan Cutler

- Related to Bug #44295: mds: MDCache.cc: 6400: FAILED ceph_assert(r == 0 || r == -2) added

#10 - 04/06/2020 08:25 PM - Nathan Cutler

Zheng, is this a follow-on fix for [#44295](#) ?

I'm asking because that issue is marked for backport all the way to mimic, but this one only to octopus.

Should we mark the mimic and nautilus backports of [#44295](#) as "Rejected"?

#11 - 04/06/2020 10:00 PM - Greg Farnum

- Backport changed from octopus to mimic, nautilus, octopus

Nathan Cutler wrote:

Zheng, is this a follow-on fix for [#44295](#) ?

I'm asking because that issue is marked for backport all the way to mimic, but this one only to octopus.

Should we mark the mimic and nautilus backports of [#44295](#) as "Rejected"?

Yes please! Also updated this ticket's backport field.

#12 - 04/10/2020 08:53 AM - Nathan Cutler

- Copied to Backport #45026: mimic: mds/Mutation.h: 128: FAILED ceph_assert(num_auth_pins == 0) added

#13 - 04/10/2020 08:54 AM - Nathan Cutler

- Copied to Backport #45027: nautilus: mds/Mutation.h: 128: FAILED ceph_assert(num_auth_pins == 0) added

#14 - 04/10/2020 08:54 AM - Nathan Cutler

- Copied to Backport #45028: octopus: mds/Mutation.h: 128: FAILED ceph_assert(num_auth_pins == 0) added

#15 - 06/03/2020 12:20 PM - Nathan Cutler

- Status changed from Pending Backport to Resolved

While running with --resolve-parent, the script "backport-create-issue" noticed that all backports of this issue are in status "Resolved" or "Rejected".