

rgw - Bug #20897

rgw: the overloaded comparison operators of rgw_bucket disrespect tenants

08/03/2017 04:37 PM - Radoslaw Zarzynski

Status:	Need Review	Start date:	08/03/2017
Priority:	Normal	Due date:	
Assignee:	Radoslaw Zarzynski	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:	Development	Reviewed:	
Tags:		Affected Versions:	
Backport:	jewel, kraken	ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor	Crash signature:	
Description			

History

#1 - 08/03/2017 04:46 PM - Radoslaw Zarzynski

- Status changed from New to Need Review

PR: <https://github.com/ceph/ceph/pull/16796>.

#2 - 08/03/2017 06:06 PM - Orit Wasserman

- Backport set to jewel, kraken

#3 - 08/04/2017 03:16 PM - Radoslaw Zarzynski

Stuffing `rgw_bucket::operator<>` with `ceph_abort()` showed it's involved in quota-related calculations:

```
ceph version 12.1.2-90-g250cfda (250cfda35819724159e585bb493476f5432c4b1e) luminous (rc)
1: (()+0x1e8944) [0x2f0944]
2: (()+0x115b0) [0x5d255b0]
3: (gsignal()+0x9f) [0x111b98df]
4: (abort()+0x16a) [0x111bb4da]
5: (()+0x563494) [0x66b494]
6: (RGWQuotaCache<rgw_bucket>::get_stats(rgw_user const&, rgw_bucket const&, RGWStorageStats&, RGWQuotaInfo&
+0xa5) [0x674f25]
7: (RGWQuotaHandlerImpl::check_bucket_shards(unsigned long, unsigned long, rgw_user const&, rgw_bucket const&
, RGWQuotaInfo&, unsigned long, bool&, unsigned int*)+0x80) [0x675250]
8: (RGWRados::check_bucket_shards(RGWBucketInfo const&, rgw_bucket const&, RGWQuotaInfo&)+0x91) [0x413861]
9: (RGWPutObj::execute()+0x1682) [0x3e13a2]
10: (rgw_process_authenticated(RGWHandler_REST*, RGWOp*, RGWRequest*, req_state*, bool)+0x172) [0x40b4c2]
11: (process_request(RGWRados*, RGWREST*, RGWRequest*, std::__cxx11::basic_string<char, std::char_traits<char
>, std::allocator<char> > const&, rgw::auth::StrategyRegistry const&, RGWRestfulIO*, OpsLogSocket*)+0x1c0c) [0
x40d4ac]
12: (RGWCivetWebFrontend::process(mg_connection*)+0x37b) [0x2aa1fb]
13: (()+0x1d95c5) [0x2e15c5]
14: (()+0x1db030) [0x2e3030]
15: (()+0x773a) [0x5d1b73a]
16: (clone()+0x3f) [0x1128be7f]
```

`RGWQuotaCache<...>::get_stats()` locates appropriate instance of `RGWQuotaCacheStats` in the LRU map. As the comparison is broken, the returned stats can be inconsistent.

```

template<class T>
int RGWQuotaCache<T>::get_stats(const rgw_user& user, const rgw_bucket& bucket, RGWStorageStats& stats, RGWQuotaInfo& quota) {
    RGWQuotaCacheStats qs;
    utime_t now = ceph_clock_now();
    if (map_find(user, bucket, qs) {
        if (qs.async_refresh_time.sec() > 0 && now >= qs.async_refresh_time) {
            int r = async_refresh(user, bucket, qs);
            if (r < 0) {
                ldout(store->ctx(), 0) << "ERROR: quota async refresh returned ret=" << r << endl;

                /* continue processing, might be a transient error, async refresh is just optimization */
            }
        }

        if (can_use_cached_stats(quota, qs.stats) && qs.expiration >
            ceph_clock_now()) {
            stats = qs.stats;
            return 0;
        }
    }

    int ret = fetch_stats_from_storage(user, bucket, stats);
    if (ret < 0 && ret != -ENOENT)
        return ret;

    set_stats(user, bucket, qs, stats);

    return 0;
}

```

Similar experiment for `rgw_bucket::operator==(())`:

```

ceph version 12.1.2-90-g250cfda (250cfda35819724159e585bb493476f5432c4b1e) luminous (rc)
1: ((()+0x1e9234) [0x2f1234]
2: ((()+0x115b0) [0x5d255b0]
3: (gsignal()+0x9f) [0x111b98df]
4: (abort()+0x16a) [0x111bb4da]
5: ((()+0x2f65f2) [0x3fe5f2]
6: (RGWRados::Object::Write::_do_write_meta(unsigned long, unsigned long, std::map<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> >, ceph::buffer::list, std::less<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > >, std::allocator<std::pair<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > const, ceph::buffer::list> > >&, bool, void*)+0xff0) [0x465610]
7: (RGWRados::Object::Write::write_meta(unsigned long, unsigned long, std::map<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> >, ceph::buffer::list, std::less<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > >, std::allocator<std::pair<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > const, ceph::buffer::list> > >&)+0x21a) [0x46585a]
8: (RGWPutObjProcessor_Atomic::do_complete(unsigned long, std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > const&, std::chrono::time_point<ceph::time_detail::real_clock, std::chrono::duration<unsigned long, std::ratio<11, 10000000001> > >*, std::chrono::time_point<ceph::time_detail::real_clock, std::chrono::duration<unsigned long, std::ratio<11, 10000000001> > >, std::map<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> >, ceph::buffer::list, std::less<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > >, std::allocator<std::pair<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > const, ceph::buffer::list> > >&, std::chrono::time_point<ceph::time_detail::real_clock, std::chrono::duration<unsigned long, std::ratio<11, 10000000001> > >, char const*, char const*, char const*, std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > const*, std::set<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> >, std::less<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > >, std::allocator<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > > >)+0x2ae) [0x465b6e]
9: (RGWPutObjProcessor::complete(unsigned long, std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > const&, std::chrono::time_point<ceph::time_detail::real_clock, std::chrono::duration<unsigned long, std::ratio<11, 10000000001> > >*, std::chrono::time_point<ceph::time_detail::real_clock, std::chrono::duration<unsigned long, std::ratio<11, 10000000001> > >, std::map<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> >, ceph::buffer::list, std::less<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > >, std::allocator<std::pair<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > const, ceph::buffer::list> > >&, std::chrono::time_point<ceph::time_detail::real_clock, std::chrono::duration<unsigned long, std::ratio<11, 10000000001> > >, char const*, char const*, std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > const*, std::set<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> >, std::less<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > >, std::allocator<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > > >, std::allocator<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > > >)+0x2ae) [0x465b6e]

```

```
raits<char>, std::allocator<char> > > >*)+0x22) [0x4123b2]
10: (RGWPutObj::execute()+0x22ed) [0x3e2eed]
11: (rgw_process_authenticated(RGWHandler_REST*, RGWOp*&, RGWRequest*, req_state*, bool)+0x172) [0x40c572]
12: (process_request(RGWRados*, RGWREST*, RGWRequest*, std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char> > const&, rgw::auth::StrategyRegistry const&, RGWRestfulIO*, OpsLogSocket*))+0x1c0c) [0x40e55c]
13: (RGWCivetWebFrontend::process(mg_connection*))+0x37b) [0x2aaaeb]
14: (()+0x1d9eb5) [0x2e1eb5]
15: (()+0x1db920) [0x2e3920]
16: (()+0x773a) [0x5d1b73a]
17: (clone()+0x3f) [0x1128be7f]
```