

## rgw - Bug #40678

### RGW S3Website didn't do the necessary checking to the website configuration

07/05/2019 03:21 PM - Enming Zhang

<b>Status:</b> Resolved	<b>% Done:</b> 0%
<b>Priority:</b> Normal	<b>Spent time:</b> 0.00 hour
<b>Assignee:</b>	
<b>Category:</b>	
<b>Target version:</b>	
<b>Source:</b>	<b>Affected Versions:</b>
<b>Tags:</b> website	<b>ceph-qa-suite:</b>
<b>Backport:</b> luminous mimic nautilus	<b>Pull request ID:</b> 28904
<b>Regression:</b> No	<b>Crash signature (v1):</b>
<b>Severity:</b> 3 - minor	<b>Crash signature (v2):</b>
<b>Reviewed:</b>	

**Description**

According to AWS S3, <https://docs.aws.amazon.com/AmazonS3/latest/API/RESTBucketPUTwebsite.html>

- To redirect all website requests sent to the bucket's website endpoint, you add a website configuration with the following elements. Because all requests are sent to another website, you don't need to provide index document name for the bucket.
- If you want granular control over redirects, you can use the following elements to add routing rules that describe conditions for redirecting requests and information about the redirect destination. In this case, the website configuration must provide an index document for the bucket, because some requests might not be redirected.

But now, the RGW don't do the necessary checking to the IndexDocument field when the RedirectAllRequestsTo is not provided, which will lead the IndexDocument Suffix to be empty without any error being reported.

In this case if we access the static website through the S3Website enabled RGW instance the RGW instance will crash.

```
In function 'void RGWObjectCtxImpl<T, S>::set_atomic(T&) [with T = rgw_obj; S = RGWObjState]' thread 7f36591ea700 time 2019-07-02 22:26:22.096930 /ceph/src/rgw/rgw_rados.h: 2164: FAILED assert(!obj.empty())
```

**Related issues:**

Duplicated by rgw - Bug #41370: [RGW] RGW in website mode: rgw_rados.h: 2150:...	<b>Duplicate</b>
Copied to rgw - Backport #41496: mimic: RGW S3Website didn't do the necessary...	<b>Resolved</b>
Copied to rgw - Backport #41497: luminous: RGW S3Website didn't do the necess...	<b>Rejected</b>
Copied to rgw - Backport #41498: nautilus: RGW S3Website didn't do the necess...	<b>Resolved</b>

### History

#1 - 07/05/2019 03:31 PM - Enming Zhang

<https://github.com/ceph/ceph/pull/28904>

#2 - 07/08/2019 05:26 PM - Casey Bodley

- Status changed from New to Fix Under Review

- Tags set to website

- Backport set to luminous mimic nautilus

- Pull request ID set to 28904

#3 - 07/10/2019 08:01 AM - Enming Zhang

The stack trace is

```
-84> 2019-07-09T23:17:54.272-0700 7f50565df700 20 HTTP_ACCEPT=**
```

```

-83> 2019-07-09T23:17:54.272-0700 7f50565df700 20 HTTP_HOST=localhost
-82> 2019-07-09T23:17:54.272-0700 7f50565df700 20 HTTP_USER_AGENT=curl/7.29.0
-81> 2019-07-09T23:17:54.272-0700 7f50565df700 20 HTTP_VERSION=1.1
-80> 2019-07-09T23:17:54.272-0700 7f50565df700 20 REMOTE_ADDR=127.0.0.1
-79> 2019-07-09T23:17:54.272-0700 7f50565df700 20 REQUEST_METHOD=GET
-78> 2019-07-09T23:17:54.272-0700 7f50565df700 20 REQUEST_URI=/webtest
-77> 2019-07-09T23:17:54.272-0700 7f50565df700 20 SCRIPT_URI=/webtest
-76> 2019-07-09T23:17:54.272-0700 7f50565df700 20 SERVER_PORT=80
-75> 2019-07-09T23:17:54.272-0700 7f50565df700 1 ===== starting new request req=0x55b572034890 =====
-74> 2019-07-09T23:17:54.272-0700 7f50565df700 1 lockdep reusing last freed id 55
-73> 2019-07-09T23:17:54.272-0700 7f50565df700 1 lockdep using id 587
-72> 2019-07-09T23:17:54.272-0700 7f50565df700 2 req 2 0.000s initializing for trans_id = tx00000000000000
0000002-005d258312-1065-default
-71> 2019-07-09T23:17:54.272-0700 7f50565df700 10 rgw api priority: s3=-1 s3website=1
-70> 2019-07-09T23:17:54.272-0700 7f50565df700 10 host=localhost
-69> 2019-07-09T23:17:54.272-0700 7f50565df700 20 subdomain= domain= in_hosted_domain=0 in_hosted_domain_s3
website=0
-68> 2019-07-09T23:17:54.273-0700 7f50565df700 -1 res_query() failed
-67> 2019-07-09T23:17:54.273-0700 7f50565df700 20 final domain/bucket subdomain= domain= in_hosted_domain=0
in_hosted_domain_s3website=1 s->info.domain= s->info.request_uri=/webtest
-66> 2019-07-09T23:17:54.273-0700 7f50565df700 20 req 2 0.001s get_handler handler=32RGWHandler_REST_Bucket
_S3Website
-65> 2019-07-09T23:17:54.273-0700 7f50565df700 10 handler=32RGWHandler_REST_Bucket_S3Website
-64> 2019-07-09T23:17:54.273-0700 7f50565df700 2 req 2 0.001s getting op 0
-63> 2019-07-09T23:17:54.273-0700 7f50565df700 10 req 2 0.001s s3:get_obj scheduling with dmclock client=2
cost=1
-62> 2019-07-09T23:17:54.273-0700 7f50565df700 10 op=28RGWGetObj_ObjStore_S3Website
-61> 2019-07-09T23:17:54.273-0700 7f50565df700 2 req 2 0.001s s3:get_obj verifying requester
-60> 2019-07-09T23:17:54.273-0700 7f50565df700 20 req 2 0.001s s3:get_obj rgw::auth::StrategyRegistry::s3_m
ain_strategy_t: trying rgw::auth::s3::AWSAuthStrategy
-59> 2019-07-09T23:17:54.273-0700 7f50565df700 20 req 2 0.001s s3:get_obj rgw::auth::s3::AWSAuthStrategy: t
rying rgw::auth::s3::S3AnonymousEngine
-58> 2019-07-09T23:17:54.273-0700 7f50565df700 20 req 2 0.001s s3:get_obj rgw::auth::s3::S3AnonymousEngine
granted access
-57> 2019-07-09T23:17:54.273-0700 7f50565df700 20 req 2 0.001s s3:get_obj rgw::auth::s3::AWSAuthStrategy gr
anted access
-56> 2019-07-09T23:17:54.273-0700 7f50565df700 2 req 2 0.001s s3:get_obj normalizing buckets and tenants
-55> 2019-07-09T23:17:54.273-0700 7f50565df700 10 s->object=<NULL> s->bucket=webtest
-54> 2019-07-09T23:17:54.273-0700 7f50565df700 2 req 2 0.001s s3:get_obj init permissions
-53> 2019-07-09T23:17:54.274-0700 7f50565df700 20 get_system_obj_state: rctx=0x55b5720336e0 obj=default.rgw
.meta:root:webtest state=0x55b570f7f840 s->prefetch_data=0
-52> 2019-07-09T23:17:54.274-0700 7f50565df700 10 cache get: name=default.rgw.meta+root+webtest : miss
-51> 2019-07-09T23:17:54.274-0700 7f50565df700 20 WARNING: blocking librados call
-50> 2019-07-09T23:17:54.274-0700 7f50565df700 1 lockdep using id 588
-49> 2019-07-09T23:17:54.275-0700 7f50565df700 10 cache put: name=default.rgw.meta+root+webtest info.flags=
0x16
-48> 2019-07-09T23:17:54.275-0700 7f50565df700 10 adding default.rgw.meta+root+webtest to cache LRU end
-47> 2019-07-09T23:17:54.275-0700 7f50565df700 10 updating xattr: name=ceph.objclass.version bl.length()=42
-46> 2019-07-09T23:17:54.275-0700 7f50565df700 20 get_system_obj_state: s->obj_tag was set empty
-45> 2019-07-09T23:17:54.275-0700 7f50565df700 20 Read xattr: user.rgw.idtag
-44> 2019-07-09T23:17:54.275-0700 7f50565df700 10 cache get: name=default.rgw.meta+root+webtest : type miss
(requested=0x13, cached=0x16)
-43> 2019-07-09T23:17:54.276-0700 7f50565df700 20 rados->read ofs=0 len=0
-42> 2019-07-09T23:17:54.276-0700 7f50565df700 20 WARNING: blocking librados call
-41> 2019-07-09T23:17:54.276-0700 7f50565df700 1 lockdep reusing last freed id 588
-40> 2019-07-09T23:17:54.277-0700 7f50565df700 20 rados_obj.operate() r=0 bl.length=171
-39> 2019-07-09T23:17:54.277-0700 7f50565df700 10 cache put: name=default.rgw.meta+root+webtest info.flags=
0x13
-38> 2019-07-09T23:17:54.277-0700 7f50565df700 10 moving default.rgw.meta+root+webtest to cache LRU end
-37> 2019-07-09T23:17:54.277-0700 7f50565df700 10 updating xattr: name=ceph.objclass.version bl.length()=42
-36> 2019-07-09T23:17:54.277-0700 7f50565df700 20 rgw_get_bucket_info: bucket instance: :webtest[e362fb6d-4
c68-4de2-aeda-c010d08ca7a3.4169.1])
-35> 2019-07-09T23:17:54.277-0700 7f50565df700 20 reading from default.rgw.meta:root:.bucket.meta.webtest:e
362fb6d-4c68-4de2-aeda-c010d08ca7a3.4169.1
-34> 2019-07-09T23:17:54.277-0700 7f50565df700 20 get_system_obj_state: rctx=0x55b5720336e0 obj=default.rgw
.meta:root:.bucket.meta.webtest:e362fb6d-4c68-4de2-aeda-c010d08ca7a3.4169.1 state=0x55b570f7fb40 s->prefetch_d
ata=0
-33> 2019-07-09T23:17:54.277-0700 7f50565df700 10 cache get: name=default.rgw.meta+root+.bucket.meta.webtes
t:e362fb6d-4c68-4de2-aeda-c010d08ca7a3.4169.1 : miss
-32> 2019-07-09T23:17:54.278-0700 7f50529d9700 10 cache put: name=default.rgw.meta+root+.bucket.meta.webtes
t:e362fb6d-4c68-4de2-aeda-c010d08ca7a3.4169.1 info.flags=0x16
-31> 2019-07-09T23:17:54.278-0700 7f50529d9700 10 adding default.rgw.meta+root+.bucket.meta.webtest:e362fb6
d-4c68-4de2-aeda-c010d08ca7a3.4169.1 to cache LRU end
-30> 2019-07-09T23:17:54.278-0700 7f50529d9700 10 updating xattr: name=ceph.objclass.version bl.length()=42
-29> 2019-07-09T23:17:54.278-0700 7f50529d9700 10 updating xattr: name=user.rgw.acl bl.length()=147

```

```

-28> 2019-07-09T23:17:54.278-0700 7f50529d9700 20 get_system_obj_state: s->obj_tag was set empty
-27> 2019-07-09T23:17:54.278-0700 7f50529d9700 20 Read xattr: user.rgw.acl
-26> 2019-07-09T23:17:54.278-0700 7f50529d9700 20 Read xattr: user.rgw.idtag
-25> 2019-07-09T23:17:54.278-0700 7f50529d9700 10 cache get: name=default.rgw.meta+root+.bucket.meta.webtest:e362fb6d-4c68-4de2-aeda-c010d08ca7a3.4169.1 : type miss (requested=0x13, cached=0x16)
-24> 2019-07-09T23:17:54.278-0700 7f50529d9700 20 rados->read ofs=0 len=0
-23> 2019-07-09T23:17:54.280-0700 7f5051fd8700 20 rados_obj.operate() r=0 bl.length=331
-22> 2019-07-09T23:17:54.280-0700 7f5051fd8700 10 cache put: name=default.rgw.meta+root+.bucket.meta.webtest:e362fb6d-4c68-4de2-aeda-c010d08ca7a3.4169.1 info.flags=0x13
-21> 2019-07-09T23:17:54.280-0700 7f5051fd8700 10 moving default.rgw.meta+root+.bucket.meta.webtest:e362fb6d-4c68-4de2-aeda-c010d08ca7a3.4169.1 to cache LRU end
-20> 2019-07-09T23:17:54.280-0700 7f5051fd8700 10 updating xattr: name=ceph.objclass.version bl.length()=42
-19> 2019-07-09T23:17:54.280-0700 7f5051fd8700 10 updating xattr: name=user.rgw.acl bl.length()=147
-18> 2019-07-09T23:17:54.280-0700 7f5051fd8700 10 chain_cache_entry: cache_locator=default.rgw.meta+root+webtest
-17> 2019-07-09T23:17:54.280-0700 7f5051fd8700 10 chain_cache_entry: cache_locator=default.rgw.meta+root+.bucket.meta.webtest:e362fb6d-4c68-4de2-aeda-c010d08ca7a3.4169.1
-16> 2019-07-09T23:17:54.280-0700 7f5051fd8700 15 decode_policy Read AccessControlPolicy<AccessControlPolicy xmlns="http://s3.amazonaws.com/doc/2006-03-01/"><Owner><ID>emtest1</ID><DisplayName>emtest1</DisplayName></Owner><AccessControlList><Grant><Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:type="CanonicalUser"><ID>emtest1</ID><DisplayName>emtest1</DisplayName></Grantee><Permission>FULL_CONTROL</Permission></Grant></AccessControlList></AccessControlPolicy>
-15> 2019-07-09T23:17:54.280-0700 7f5051fd8700 20 get_system_obj_state: rctx=0x55b572033050 obj=default.rgw.meta:users.uid:anonymous state=0x55b570f7fe40 s->prefetch_data=0
-14> 2019-07-09T23:17:54.280-0700 7f5051fd8700 10 cache get: name=default.rgw.meta+users.uid+anonymous : type miss (requested=0x6, cached=0x0)
-13> 2019-07-09T23:17:54.280-0700 7f5051fd8700 20 WARNING: blocking librados call
-12> 2019-07-09T23:17:54.280-0700 7f5051fd8700 1 lockdep reusing last freed id 588
-11> 2019-07-09T23:17:54.281-0700 7f5051fd8700 10 cache put: name=default.rgw.meta+users.uid+anonymous info.flags=0x0
-10> 2019-07-09T23:17:54.281-0700 7f5051fd8700 10 moving default.rgw.meta+users.uid+anonymous to cache LRU end
-9> 2019-07-09T23:17:54.281-0700 7f5051fd8700 2 req 2 0.009s s3:get_obj recalculating target
-8> 2019-07-09T23:17:54.281-0700 7f5051fd8700 10 req 2 0.009s retarget Starting retarget
-7> 2019-07-09T23:17:54.281-0700 7f5051fd8700 10 req 2 0.009s retarget get_effective_key ->
-6> 2019-07-09T23:17:54.281-0700 7f5051fd8700 2 req 2 0.009s s3:get_obj reading permissions
-5> 2019-07-09T23:17:54.281-0700 7f5051fd8700 2 req 2 0.009s s3:get_obj init op
-4> 2019-07-09T23:17:54.281-0700 7f5051fd8700 2 req 2 0.009s s3:get_obj verifying op mask
-3> 2019-07-09T23:17:54.281-0700 7f5051fd8700 20 req 2 0.009s s3:get_obj required_mask= 1 user.op_mask=7
-2> 2019-07-09T23:17:54.281-0700 7f5051fd8700 2 req 2 0.009s s3:get_obj verifying op permissions
-1> 2019-07-09T23:17:54.281-0700 7f5051fd8700 0 [debug by vampirem] bucket :webtest[e362fb6d-4c68-4de2-aeda-c010d08ca7a3.4169.1]) object
0> 2019-07-09T23:17:54.296-0700 7f5051fd8700 -1 *** Caught signal (Aborted) **
in thread 7f5051fd8700 thread_name:radosgw

```

```

ceph version v15.0.0-2498-g9b967b4 (9b967b42f09f435cfabd0937f6a5c5912379d9a6) octopus (dev)
1: ((()+0xd8d974) [0x55b56e724974]
2: ((()+0xf5d0) [0x7f508af3b5d0]
3: (gsignal()+0x37) [0x7f5089ca2207]
4: (abort()+0x148) [0x7f5089ca38f8]
5: ((()+0x2f026) [0x7f5089c9b026]
6: ((()+0x2f0d2) [0x7f5089c9b0d2]
7: (RGWObjectCtx::set_atomic(rgw_obj&)+0x69) [0x55b56e7b50ad]
8: (RGWRados::set_atomic(void*, rgw_obj&)+0x2f) [0x55b56eb37d65]
9: (RGWGetObj::verify_permission()+0x1e2) [0x55b56eaf314c]
10: (rgw_process_authenticated(RGWHandler_REST*, RGWOp*&, RGWRequest*, req_state*, bool)+0xd1e) [0x55b56e76681f]
11: (process_request(RGWRados*, RGWREST*, RGWRequest*, std::string const&, rgw::auth::StrategyRegistry const&, RGWRestfulIO*, OpsLogSocket*, optional_yield, rgw::dmclock::Scheduler*, int*)+0x1483) [0x55b56e768d47]
12: ((()+0xc81cf6) [0x55b56e618cf6]
13: ((()+0xc7d81a) [0x55b56e61481a]
14: ((()+0xc8949c) [0x55b56e62049c]
15: ((()+0xc8915a) [0x55b56e62015a]
16: ((()+0xc88e80) [0x55b56e61fe80]
17: (make_fcontext()+0x2f) [0x55b56ede6b2f]
NOTE: a copy of the executable, or `objdump -rdS <executable>` is needed to interpret this.

```

**#4 - 08/16/2019 02:26 PM - Casey Bodley**

- Status changed from Fix Under Review to Pending Backport

**#5 - 08/26/2019 02:55 PM - Nathan Cutler**

- Copied to Backport #41496: mimic: RGW S3Website didn't do the necessary checking to the website configuration added

**#6 - 08/26/2019 02:55 PM - Nathan Cutler**

- Copied to Backport #41497: luminous: RGW S3Website didn't do the necessary checking to the website configuration added

**#7 - 08/26/2019 02:56 PM - Nathan Cutler**

- Copied to Backport #41498: nautilus: RGW S3Website didn't do the necessary checking to the website configuration added

**#8 - 10/17/2019 04:51 PM - Nathan Cutler**

- Project changed from Ceph to rgw

**#9 - 08/18/2020 10:22 PM - Nathan Cutler**

- Duplicated by Bug #41370: [RGW] RGW in website mode: rgw\_rados.h: 2150: FAILED assert(!obj.empty()) added

**#10 - 01/27/2021 08:06 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".