

## rgw - Bug #43756

### An error occurred (NoSuchUpload) when calling the AbortMultipartUpload operation: Unknown

01/22/2020 03:16 PM - Manuel Rios

<b>Status:</b>	Triaged	<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>	v14.2.6
<b>Tags:</b>		<b>ceph-qa-suite:</b>	
<b>Backport:</b>		<b>Pull request ID:</b>	
<b>Regression:</b>	No	<b>Crash signature (v1):</b>	
<b>Severity:</b>	2 - major	<b>Crash signature (v2):</b>	
<b>Reviewed:</b>			

#### Description

Hi RGW Team,

The last 7 days we spent trying to solve a metering problem in the buckets.

Well, right now its looks like LifeCycle are not able to purge/delete some objects may be due to some parse problem.

Let post some information:

```
radosgw-admin user stats --uid=XXXXX
{
  "stats": {
    "total_entries": 22817077,
    "total_bytes": 164278075532090,
    "total_bytes_rounded": 164325122670592
  },
  "last_stats_sync": "2020-01-21 19:32:02.231796Z",
  "last_stats_update": "2020-01-22 14:48:30.915696Z"
}
```

Aprox 164TB usage.

The customer got near 57 buckets with different sizes I'm going to post just one.

```
radosgw-admin bucket stats --bucket=Evol6
{
  "bucket": "Evol6",
  "tenant": "",
  "zonegroup": "4d8c7c5f-ca40-4ee3-b5bb-b2cad90bd007",
  "placement_rule": "default-placement",
  "explicit_placement": {
    "data_pool": "default.rgw.buckets.data",
    "data_extra_pool": "default.rgw.buckets.non-ec",
    "index_pool": "default.rgw.buckets.index"
  },
  "id": "48efb8c3-693c-4fe0-bbe4-fdc16f590a82.132873679.2",
  "marker": "48efb8c3-693c-4fe0-bbe4-fdc16f590a82.3886182.52",
  "index_type": "Normal",
  "owner": "xxxxxx",
  "ver": "0#91266,1#60635,2#80715,3#78528",
```

```

"master_ver": "0#0,1#0,2#0,3#0",
"mtime": "2020-01-21 22:38:31.437616Z",
"max_marker": "0#,1#,2#,3#",
"usage": {
  "rgw.main": {
    "size": 9107173119747,
    "size_actual": 9107345551360,
    "size_utilized": 9107173119747,
    "size_kb": 8893723750,
    "size_kb_actual": 8893892140,
    "size_kb_utilized": 8893723750,
    "num_objects": 180808
  },
  "rgw.multimeta": {
    "size": 0,
    "size_actual": 0,
    "size_utilized": 3807,
    "size_kb": 0,
    "size_kb_actual": 0,
    "size_kb_utilized": 4,
    "num_objects": 141
  }
},
"bucket_quota": {
  "enabled": false,
  "check_on_raw": false,
  "max_size": -1024,
  "max_size_kb": 0,
  "max_objects": -1
}
}

```

Current size 9 TB approx. external Tools like S3 Browser // Cloudberry Amazon S3 Explorer reports 7TB.  
Its a considerable difference but its not metadata overhead.

Checking with AWS CLI we found incomplected multipart, it's normal due customer backup thousand of remote computers and use CEPH as backend.

I found a small script to cancel all multipart using AWS CLI.

```

BUCKETNAME=Evol6
aws --endpoint=http://XXXXXX:7480 --profile=ceph s3api list-multipart-uploads --bucket $BUCKETNAME \
> | jq -r '.Uploads[] | "--key \"(.Key)\" --upload-id \"(.UploadId)\"' \
> | while read -r line; do
>   eval "aws --endpoint=http://XXXXXXX:7480 --profile=ceph s3api abort-multipart-upload --bucket $BUCKETNAME $line";
> done

```

The output generates the same error:

```
An error occurred (NoSuchUpload) when calling the AbortMultipartUpload operation: Unknown
```

```
An error occurred (NoSuchUpload) when calling the AbortMultipartUpload operation: Unknown
```

```
An error occurred (NoSuchUpload) when calling the AbortMultipartUpload operation: Unknown
```

Checking the multipart list :

```

{
  "Initiator": {
    "DisplayName": "xxxxx",
    "ID": "xxxxx"
  },
  "Initiated": "2019-12-03T02:00:50.589Z",
  "UploadId": "2~T7G76R09Pn-267VMbY8cJvZ1_BHqfTx",
  "StorageClass": "STANDARD",
  "Key": "MBS-da43656f-2b8c-464f-b341-03fdbdf446ae/CBB_SRV2K12/CBB_VM/192.168.0.197/SRV2
K12/Hard disk 1$/20191203010516/431.cbrevision",
  "Owner": {
    "DisplayName": "xxxxx",
    "ID": "xxxxx"
  }
},
{
  "Initiator": {
    "DisplayName": "xxxxx",
    "ID": "xxxxx"
  },
  "Initiated": "2019-12-03T01:23:06.007Z",
  "UploadId": "2~r0BMPPs8CewVZ6Qheuls9WzaBn7bBvU",
  "StorageClass": "STANDARD",
  "Key": "MBS-da43656f-2b8c-464f-b341-03fdbdf446ae/CBB_SRV2K12/CBB_VM/192.168.0.197/SRV2
K12/Hard disk 1$/20191203010516/431.cbrevision",
  "Owner": {
    "DisplayName": "xxxxx",
    "ID": "xxxxx"
  }
}
}

```

Maybe the parse internally of 1\$ char is generating a problem in the LC scripts that don't allow getting purged.

The main problem of that issue is the huge difference between completed files that show in all external tools and the internal storage metering.

Additionally for help this type of customer we add a LC policy that for some reason fails but shows as completed.

```

s3cmd getlifecycle s3://Evol6 --no-ssl
<?xml version="1.0" ?>
<LifecycleConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
    <ID>Incomplete Multipart Uploads</ID>
    <Prefix/>
    <Status>Enabled</Status>
    <AbortIncompleteMultipartUpload>
      <DaysAfterInitiation>1</DaysAfterInitiation>
    </AbortIncompleteMultipartUpload>
  </Rule>
</LifecycleConfiguration>

```

```

radosgw-admin lc list
[
  {
    "bucket": ":Evol6:48efb8c3-693c-4fe0-bbe4-fdc16f590a82.3886182.52",
    "status": "COMPLETE"
  }
]

```

Obviously completed is not the correct error because in the multipart incomplete show like 157 incompleted uploads.

I appreciated all the help and ideas.

Best Regards

**Related issues:**

Related to rgw - Bug #43583: rgw: unable to abort multipart upload after the ...

**Resolved**

**History**

**#1 - 01/22/2020 05:07 PM - Manuel Rios**

Hi

With help we launched a standalone rgw instance with a non public port and launched just 3 commands with AWS CLI

```
aws --endpoint=http://XXXXXX:7481 --profile=ceph s3api list-multipart-uploads --bucket Evol6
```

```
aws --endpoint=http://XXXXXX:7481 --profile=ceph s3api list-parts --bucket Evol6 --key 'MBS-da43656f-2b8c-464f-b341-03fdbdf446ae/CBB_SRV2K12/CBB_VM/192.168.0.197/SRV2K12/Hard disk 1$/20191203010516/431.cbrevision' --upload-id 2~r0BMPPs8CewVZ6Qheuls9WzaBn7bBvU
```

```
aws --endpoint=http://XXXXXX:7481 --profile=ceph s3api abort-multipart-upload --bucket Evol6 --key 'MBS-da43656f-2b8c-464f-b341-03fdbdf446ae/CBB_SRV2K12/CBB_VM/192.168.0.197/SRV2K12/Hard disk 1$/20191203010516/431.cbrevision' --upload-id 2~r0BMPPs8CewVZ6Qheuls9WzaBn7bBvU
```

RGW log output with <https://easydatahost.com/debugs/debug-rgw.zip>

RGW DAEMON :

```
/usr/bin/radosgw -d --cluster ceph --name client.rgw.ceph-rgw03 --setuser ceph --setgroup ceph --debug-rgw=20 --debug_ms=1 --rgw_frontends="beast port=7481" --rgw_enable_gc_threads=false --rgw_enable_lc_threads=false
```

**#2 - 01/22/2020 05:22 PM - Manuel Rios**

Output of cli : radosgw-admin bi list --bucket Evol6 | jq '[]|select(.idx | match("20191203010516/431.cbrevision"))'

```
{
  "type": "plain",
  "idx": "_multipart_MBS-da43656f-2b8c-464f-b341-03fdbdf446ae/CBB_SRV2K12/CBB_VM/192.168.0.197/SRV2K12/Hard disk 1$/20191203010516/431.cbrevision.2~T7G76R09Pn-267VMbY8cjbvZl_BHqfTx.meta",
  "entry": {
    "name": "_multipart_MBS-da43656f-2b8c-464f-b341-03fdbdf446ae/CBB_SRV2K12/CBB_VM/192.168.0.197/SRV2K12/Hard disk 1$/20191203010516/431.cbrevision.2~T7G76R09Pn-267VMbY8cjbvZl_BHqfTx.meta",
    "instance": "",
    "ver": {
      "pool": 40,
```

```

    "epoch": 4848481
  },
  "locator": "",
  "exists": "true",
  "meta": {
    "category": 3,
    "size": 27,
    "mtime": "2019-12-03 02:00:50.589889Z",
    "etag": "",
    "storage_class": "",
    "owner": "catbackup",
    "owner_display_name": "Catbackup",
    "content_type": "application/octet-stream",
    "accounted_size": 0,
    "user_data": "",
    "appendable": "false"
  },
  "tag": "_OQRXmFYGxL4JorOtTIVTgaWPP4Hciiu",
  "flags": 0,
  "pending_map": [],
  "versioned_epoch": 0
}
}
{
  "type": "plain",
  "idx": "_multipart_MBS-da43656f-2b8c-464f-b341-03fdbdf446ae/CBB_SRV2K12/CBB_VM/192.168.0.197/SRV2K12/Hard disk 1$/20191203010516/431.cbrevision.2~r0BMPPs8CewVZ6Qheuls9WzaBn7bBvU.meta",
  "entry": {
    "name": "_multipart_MBS-da43656f-2b8c-464f-b341-03fdbdf446ae/CBB_SRV2K12/CBB_VM/192.168.0.197/SRV2K12/Hard disk 1$/20191203010516/431.cbrevision.2~r0BMPPs8CewVZ6Qheuls9WzaBn7bBvU.meta",
    "instance": "",
    "ver": {
      "pool": 40,
      "epoch": 4862265
    },
    "locator": "",
    "exists": "true",
    "meta": {
      "category": 3,
      "size": 27,
      "mtime": "2019-12-03 01:23:06.007727Z",
      "etag": "",
      "storage_class": "",
      "owner": "catbackup",
      "owner_display_name": "Catbackup",
      "content_type": "application/octet-stream",
      "accounted_size": 0,
      "user_data": "",
      "appendable": "false"
    },
    "tag": "_ShAUoEzV6fSf9M5DGRAfIUlnlN-bCwR4",
    "flags": 0,
    "pending_map": [],
    "versioned_epoch": 0
  }
}
}
{
  "type": "plain",
  "idx": "_multipart_MBS-da43656f-2b8c-464f-b341-03fdbdf446ae/CBB_SRV2K12/CBB_VM/192.168.0.197/SRV2K12/Hard disk 1$/20191203010516/431.cbrevision.2~9djvntf2OBzWT8VlMBixPjZMx6rSwI_.meta",
  "entry": {
    "name": "_multipart_MBS-da43656f-2b8c-464f-b341-03fdbdf446ae/CBB_SRV2K12/CBB_VM/192.168.0.197/SRV2K12/Hard disk 1$/20191203010516/431.cbrevision.2~9djvntf2OBzWT8VlMBixPjZMx6rSwI_.meta",
    "instance": "",
    "ver": {
      "pool": 40,
      "epoch": 4848897
    },
    "locator": "",
    "exists": "true",
    "meta": {
      "category": 3,
      "size": 27,
      "mtime": "2019-12-03 03:00:19.076330Z",
      "etag": "",

```

```

    "storage_class": "",
    "owner": "catbackup",
    "owner_display_name": "Catbackup",
    "content_type": "application/octet-stream",
    "accounted_size": 0,
    "user_data": "",
    "appendable": "false"
  },
  "tag": "_dj5cX7yiIK3HxrLtWYollihSdkERdtL",
  "flags": 0,
  "pending_map": [],
  "versioned_epoch": 0
}
}

```

### #3 - 01/22/2020 06:01 PM - Robin Johnson

cbodley:

I sat down and debugging this with mrf.

There's a few things here, generally related:

#### 1. MPU Heads

1.1. MPU heads that are still in the index, but the .meta RADOS object is gone.

#### 2. MPU Parts

2.1. MPU parts that are still in the index but NOT RADOS, but the MPU head is missing in the index

2.2. MPU parts that are still in the index AND RADOS, but the MPU head is missing in the index

I think there was a issue for a generalized MPU cleanup tooling, but I don't know the ticket number. This shows the immediate need for it. The leaked parts are eating ~2TB of storage in just this one bucket. DigitalOcean has seem the same issue as far back as Luminous.

### #4 - 01/22/2020 06:06 PM - Robin Johnson

Snippet of logs showing the MPU head without the RADOS object:

```

2020-01-22 17:45:06.358 7f197fc31700 2 req 2 0.002s s3:list_multipart recalculating target
2020-01-22 17:45:06.358 7f197fc31700 2 req 2 0.002s s3:list_multipart reading permissions
2020-01-22 17:45:06.358 7f197fc31700 20 get_obj_state: rctx=0x564250a2c0d0 obj=Evol6:_multipart_MBS-da43656f-2
b8c-464f-b341-03fdbdf446ae/CBB_SRV2K12/CBB_VM/192.168.0.197/SRV2K12/Hard disk 1$/20191203010516/431.cbrevision
.2~r0BMPPs8CewVZ6Qheuls9WzaBn7bBvU.meta state=0x5642501820a0 s->prefetch_data=0
2020-01-22 17:45:06.358 7f197fc31700 1 -- 172.16.2.8:0/218001572 --> [v2:172.16.2.12:6852/524389,v1:172.16.2.
12:6853/524389] -- osd_op(unknown.0.0:517 40.1 40:9d5d3eed:::48efb8c3-693c-4fe0-bbe4-fdc16f590a82.3886182.52__
multipart_MBS-da43656f-2b8c-464f-b341-03fdbdf446ae%2fCBB_SRV2K12%2fCBB_VM%2f192.168.0.197%2fSRV2K12%2fHard dis
k 1$%2f20191203010516%2f431.cbrevision.2~r0BMPPs8CewVZ6Qheuls9WzaBn7bBvU.meta:head [getxattrs,stat] snapc 0=[]
ondisk+read+known_if_redirected e1097995) v8 -- 0x56424ffeedc0 con 0x56424fcf8800
2020-01-22 17:45:06.359 7f19a5c7d700 1 -- 172.16.2.8:0/218001572 <== osd.73 v2:172.16.2.12:6852/524389 10 ==
= osd_op_reply(517 48efb8c3-693c-4fe0-bbe4-fdc16f590a82.3886182.52__multipart_MBS-da43656f-2b8c-464f-b341-03fd
bdf446ae/CBB_SRV2K12/CBB_VM/192.168.0.197/SRV2K12/Hard disk 1$/20191203010516/431.cbrevision.2~r0BMPPs8CewVZ6Q
heuls9WzaBn7bBvU.meta [getxattrs,stat] v0'0 uv0 ondisk = -2 ((2) No such file or directory)) v8 ==== 408+0+0 (

```

```
crc 0 0 0) 0x56425037a280 con 0x56424fcf8800
2020-01-22 17:45:06.359 7f197fc31700 15 decode_policy Read AccessControlPolicy<AccessControlPolicy xmlns="http
://s3.amazonaws.com/doc/2006-03-01/"><Owner><ID>[SENSITIVE DATA]</ID><DisplayName>Catbackup</DisplayName></Own
er><AccessControlList><Grant><Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:type="Canonica
lUser"><ID>[SENSITIVE DATA]</ID><DisplayName>Catbackup</DisplayName></Grantee><Permission>FULL_CONTROL</Permis
sion></Grant></AccessControlList></AccessControlPolicy>
2020-01-22 17:45:06.359 7f197fc31700 10 req 2 0.003s s3:list_multipart read_permissions on Evol6[48efb8c3-693c
-4fe0-bbe4-fdc16f590a82.3886182.52]:MBS-da43656f-2b8c-464f-b341-03fdbdf446ae/CBB_SRV2K12/CBB_VM/192.168.0.197/
SRV2K12/Hard disk 1$/20191203010516/431.cbrevision only_bucket=0 ret=-2
2020-01-22 17:45:06.359 7f197fc31700 20 op->ERRORHANDLER: err_no=-2 new_err_no=-2
2020-01-22 17:45:06.359 7f197fc31700 2 req 2 0.003s s3:list_multipart op status=0
2020-01-22 17:45:06.359 7f197fc31700 2 req 2 0.003s s3:list_multipart http status=404
```

**#5 - 01/22/2020 09:36 PM - J. Eric Ivancich**

I wonder if this is affected by the bug in this tracker/pr:

- <https://tracker.ceph.com/issues/43583>
- <https://github.com/ceph/ceph/pull/32617>

Resharding wasn't putting the MPU parts on the right shards. So the question is if there has been a reshard since the multipart uploads were initiated?

Even once that PR has merged, the MPU parts would still be on the wrong shards, but a reshard would get them on the right shards.

cbodley, in an off-line discussion, suggested a possible work-around that could be done before the PR merges:

1. reshard down to ONE single shard (i.e., then everything is inherently on the right shard)
2. clean up the incomplete multipart uploads
3. reshard to the desired number of shards

I don't know that that process has been tested. If one were to test it, it might be worth trying that on a single, small-sized bucket.

**#6 - 01/22/2020 11:21 PM - Manuel Rios**

Hi Eric / Team,

Im going to test your theory about rehard to single shard -> cleanup -> reshard to XX shards.

Im going to do with a bucket of a terminated project that got the same error.

Will report results in a couple hours.

Regards

## #7 - 01/22/2020 11:42 PM - Manuel Rios

Well i got the result and nothing sucessfully :

Bucket = DMS

Multiparts incompleted date 20190912

```
<pre><code class="text">
```

```
{
  "Initiator": {
    "DisplayName": "xxxxx",
    "ID": "xxxxx"
  },
  "Initiated": "2019-09-12T01:38:03.921Z",
  "UploadId": "2~Ge19DNi2OVDTu0fqZ7fgJlH2CrIttJ",
  "StorageClass": "STANDARD",
  "Key": "MBS-8a3218ee-24a4-42aa-8535-fda31eb46a0d/CBB_MENENDEZ-TS/C$/copias_sql/Kmaleon 20190911 2230.sql$/20190911203244/Kmaleon 20190911 2230.sql",
  "Owner": {
    "DisplayName": "xxxxx",
    "ID": "xxxxx"
  }
},
{
  "Initiator": {
    "DisplayName": "xxxxx",
    "ID": "xxxx"
  },
  "Initiated": "2019-09-11T22:22:55.136Z",
  "UploadId": "2~zZOUOY1ewrhTH9CPfURkImjusiffzkt",
  "StorageClass": "STANDARD",
  "Key": "MBS-8a3218ee-24a4-42aa-8535-fda31eb46a0d/CBB_MENENDEZ-TS/C$/copias_sql/Kmaleon 20190911 2230.sql$/20190911203244/Kmaleon 20190911 2230.sql",
  "Owner": {
    "DisplayName": "xxxxxxx",
    "ID": "xxxxxxx"
  }
}
]
}
</pre>
```

```
radosgw-admin reshard add --bucket DMS --num-shards 1 --yes-i-really-mean-it
```

```
[
  {
    "time": "2020-01-22 23:22:32.807698Z",
    "tenant": "",
    "bucket_name": "DMS",
    "bucket_id": "48efb8c3-693c-4fe0-bbe4-fdc16f590a82.130777415.4",
    "new_instance_id": "",
    "old_num_shards": 32,
    "new_num_shards": 1
  }
]
```

**radosgw-admin reshard process**

```
2020-01-23 00:28:07.871 7f8edaeld6c0 1 execute INFO: reshard of bucket "DMS" from "DMS:48efb8c3-693c-4fe0-bbe4-fdc16f590a82.130777415.4" to "DMS:48efb8c3-693c-4fe0-bbe4-fdc16f590a82.134292855.1" completed successfully
```

**Checking new bucket sharding:**

```
[root@ceph-rgw03 ~]# radosgw-admin bucket stats --bucket DMS
{
```



```

"bucket": "DMS",
"tenant": "",
"zonegroup": "4d8c7c5f-ca40-4ee3-b5bb-b2cad90bd007",
"placement_rule": "default-placement",
"explicit_placement": {
  "data_pool": "",
  "data_extra_pool": "",
  "index_pool": ""
},
"id": "48efb8c3-693c-4fe0-bbe4-fdc16f590a82.134292855.1",
"marker": "48efb8c3-693c-4fe0-bbe4-fdc16f590a82.110976409.2",
"index_type": "Normal",
"owner": "xxxxxxx",
"ver": "0#28295",
"master_ver": "0#0",
"mtime": "2020-01-22 23:23:26.670489Z",
"max_marker": "0#",
"usage": {
  "rgw.main": {
    "size": 1566138931393,
    "size_actual": 1569989439488,
    "size_utilized": 1566138931393,
    "size_kb": 1529432551,
    "size_kb_actual": 1533192812,
    "size_kb_utilized": 1529432551,
    "num_objects": 1810738
  },
  "rgw.multimeta": {
    "size": 0,
    "size_actual": 0,
    "size_utilized": 459,
    "size_kb": 0,
    "size_kb_actual": 0,
    "size_kb_utilized": 1,
    "num_objects": 17
  }
},
"bucket_quota": {
  "enabled": false,
  "check_on_raw": false,
  "max_size": -1,
  "max_size_kb": 0,
  "max_objects": -1
}
}
}

```

### Count the number of multipart

```

aws --endpoint=http://xxxxxxxxx:7480 --profile=ceph s3api list-multipart-uploads --bucket $BUCKETNAME \
| jq -r '.Uploads[] | "--key \"\"(.Key)\" --upload-id \"\"(.UploadId)\"' | wc

```

Reports 17 multipart.

Now trying again the delete:

```

BUCKETNAME=DMS
aws --endpoint=http://xxxxxxxxx:7480 --profile=ceph s3api list-multipart-uploads --bucket $BUCKETNAME \
| jq -r '.Uploads[] | "--key \"\"(.Key)\" --upload-id \"\"(.UploadId)\"' \
| while read -r line; do
  eval "aws --endpoint=http://xxxxxxxxx:7480 --profile=ceph s3api abort-multipart-upload --bucket $BUCKETNAM
E $line";
done

```

```

x17 An error occurred (NoSuchUpload) when calling the AbortMultipartUpload operation: Unknown

```

**#8 - 01/22/2020 11:45 PM - Manuel Rios**

Just a note:

Once finish the reshard , bucket stats dont show anymore :

```
"explicit_placement": {  
  "data_pool": "",  
  "data_extra_pool": "",  
  "index_pool": ""  
},
```

Compared other buckets

```
"explicit_placement": {  
  "data_pool": "default.rgw.buckets.data",  
  "data_extra_pool": "default.rgw.buckets.non-ec",  
  "index_pool": "default.rgw.buckets.index"  
},
```

**#9 - 01/27/2020 05:30 PM - Manuel Rios**

Any update or workarround from developers area?

**#10 - 02/02/2020 10:43 AM - Or Friedmann**

I saw that you have / in your object name have you tried to use \ as an escape character?

I would be happy to see only the output of the abortmultipart request in the rgw log (debug-ms=0 debug-rgw=20)

Thank you

**#11 - 02/06/2020 03:08 PM - Casey Bodley**

- Status changed from New to Triaged

**#12 - 02/11/2020 07:54 AM - Manuel Rios**

Hi Mr Friedmann,

Here you can download the debug requested:

<https://file.io/u83gRj>

Regards

**#13 - 02/13/2020 03:18 PM - Casey Bodley**

- Related to Bug #43583: *rgw: unable to abort multipart upload after the bucket got resharded added*

**#14 - 03/26/2020 07:47 PM - Chris Jones**

Just an FYI... I know Jewel is EOL, but I am seeing unabortable multipart uploads in Jewel due to bucket resharding as well.