

## RADOS - Bug #40634

### mon: auth mon isn't loading full KeyServerData after restart

07/02/2019 10:59 PM - Sage Weil

<b>Status:</b>	Resolved	<b>% Done:</b>	0%
<b>Priority:</b>	Urgent	<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>		<b>ceph-qa-suite:</b>	
<b>Backport:</b>	nautilus, mimic, luminous	<b>Component(RADOS):</b>	
<b>Regression:</b>	No	<b>Pull request ID:</b>	28850
<b>Severity:</b>	3 - minor	<b>Crash signature (v1):</b>	
<b>Reviewed:</b>		<b>Crash signature (v2):</b>	

#### Description

/a/sage-2019-07-02\_17:58:21-rados-wip-sage-testing-2019-07-02-1056-distro-basic-smithi/4087648

symptom is a failed auth attempt. in this case the osd failed to start with

```
2019-07-02T18:51:31.751+0000 7f1ec28a4700 1 --2- [v2:172.21.15.1:6800/14298,v1:172.21.15.1:6801/14298] >> [v2:172.21.15.1:3302/0,v1:172.21.15.1:6791/0] conn(0x558361922000 0x558361243700 unknown : -1 s=AUTH_CONNECTING pgs=0 cs=0 l=1 rx=0 tx=0).handle_auth_bad_method method=2 result (5) Input/output error, allowed methods=[2], allowed modes=[2,1]
2019-07-02T18:51:31.752+0000 7f1ec28a4700 -1 monclient(hunting): handle_auth_bad_method server allowed_methods [2] but i only support [2]
```

on mon.c, that's because

```
2019-07-02T18:51:31.751+0000 7fef90995700 10 cephx server osd.0: handle_request auth ticket global_id 10020
2019-07-02T18:51:31.751+0000 7fef90995700 10 cephx keyserverdata: get_service_secret service auth not found
```

this is a short time after mon.c restarted, and didn't seem to load the full KeyServerData...

```
2019-07-02T18:51:29.090+0000 7fef9399b700 10 mon.c@2(peon).auth v4 update_from_paxos version 4 key server 0 latest 0
2019-07-02T18:51:29.090+0000 7fef9399b700 10 mon.c@2(peon).auth v4 update_from_paxos key server version 0
2019-07-02T18:51:29.091+0000 7fef9399b700 20 mon.c@2(peon).auth v4 update_from_paxos walking through version 1 len 120
2019-07-02T18:51:29.091+0000 7fef9399b700 20 mon.c@2(peon).auth v4 update_from_paxos walking through version 2 len 120
2019-07-02T18:51:29.091+0000 7fef9399b700 20 mon.c@2(peon).auth v4 update_from_paxos walking through version 3 len 120
2019-07-02T18:51:29.091+0000 7fef9399b700 20 mon.c@2(peon).auth v4 update_from_paxos walking through version 4 len 988
2019-07-02T18:51:29.091+0000 7fef9399b700 10 mon.c@2(peon).auth v4 update_from_paxos last_allocated_id initialized to 0
```

```
2019-07-02T18:51:29.091+0000 7fef9399b700 10 mon.c@2(peon).auth v4 update_from_paxos max_global_id=0 format_version 0
```

except v1 should have had all of the rotating keys:

```
2019-07-02T18:49:22.755+0000 7fee11643700 10 cephx keyserver: _check_rotating_secrets
2019-07-02T18:49:22.755+0000 7fee11643700 10 cephx keyserver: _rotate_secret adding auth
2019-07-02T18:49:22.755+0000 7fee11643700 10 cephx keyserver: _rotate_secret adding auth
2019-07-02T18:49:22.755+0000 7fee11643700 10 cephx keyserver: _rotate_secret adding auth
2019-07-02T18:49:22.755+0000 7fee11643700 10 cephx keyserver: _rotate_secret adding mon
2019-07-02T18:49:22.755+0000 7fee11643700 10 cephx keyserver: _rotate_secret adding mon
2019-07-02T18:49:22.755+0000 7fee11643700 10 cephx keyserver: _rotate_secret adding mon
2019-07-02T18:49:22.755+0000 7fee11643700 10 cephx keyserver: _rotate_secret adding osd
2019-07-02T18:49:22.755+0000 7fee11643700 10 cephx keyserver: _rotate_secret adding osd
2019-07-02T18:49:22.755+0000 7fee11643700 10 cephx keyserver: _rotate_secret adding osd
2019-07-02T18:49:22.755+0000 7fee11643700 10 cephx keyserver: _rotate_secret adding mds
2019-07-02T18:49:22.755+0000 7fee11643700 10 cephx keyserver: _rotate_secret adding mds
2019-07-02T18:49:22.755+0000 7fee11643700 10 cephx keyserver: _rotate_secret adding mds
2019-07-02T18:49:22.755+0000 7fee11643700 10 cephx keyserver: _rotate_secret adding mgr
2019-07-02T18:49:22.755+0000 7fee11643700 10 cephx keyserver: _rotate_secret adding mgr
2019-07-02T18:49:22.755+0000 7fee11643700 10 cephx keyserver: _rotate_secret adding mgr
2019-07-02T18:49:22.755+0000 7fee11643700 10 cephx keyserver: _check_rotating_secrets added 15
2019-07-02T18:49:22.755+0000 7fee11643700 10 mon.a@0(leader).auth v0 check_rotate updated rotating
2019-07-02T18:49:22.755+0000 7fee11643700 10 mon.a@0(leader).auth v0 get_initial_keyring
2019-07-02T18:49:22.756+0000 7fee11643700 10 mon.a@0(leader).auth v0 create_initial_keys with keyring
2019-07-02T18:49:22.756+0000 7fee11643700 10 mon.a@0(leader).auth v0 import_keyring 12 keys
2019-07-02T18:49:22.756+0000 7fee11643700 10 mon.a@0(leader).auth v0 add auth entity osd.0
2019-07-02T18:49:22.756+0000 7fee11643700 10 mon.a@0(leader).auth v0 add auth entity osd.1
2019-07-02T18:49:22.756+0000 7fee11643700 10 mon.a@0(leader).auth v0 add auth entity osd.2
2019-07-02T18:49:22.756+0000 7fee11643700 10 mon.a@0(leader).auth v0 add auth entity client.0
2019-07-02T18:49:22.756+0000 7fee11643700 10 mon.a@0(leader).auth v0 add auth entity client.admin
2019-07-02T18:49:22.756+0000 7fee11643700 10 mon.a@0(leader).auth v0 add auth entity client.boots
trap-mds
2019-07-02T18:49:22.756+0000 7fee11643700 10 mon.a@0(leader).auth v0 add auth entity client.boots
trap-mgr
2019-07-02T18:49:22.756+0000 7fee11643700 10 mon.a@0(leader).auth v0 add auth entity client.boots
trap-osd
2019-07-02T18:49:22.756+0000 7fee11643700 10 mon.a@0(leader).auth v0 add auth entity client.boots
trap-rbd
2019-07-02T18:49:22.756+0000 7fee11643700 10 mon.a@0(leader).auth v0 add auth entity client.boots
trap-rbd-mirror
2019-07-02T18:49:22.756+0000 7fee11643700 10 mon.a@0(leader).auth v0 add auth entity client.boots
trap-rgw
2019-07-02T18:49:22.756+0000 7fee11643700 10 mon.a@0(leader).auth v0 add auth entity mgr.x
2019-07-02T18:49:22.756+0000 7fee11643700 10 mon.a@0(leader).paxoservice(auth 0..0) propose_pending
2019-07-02T18:49:22.756+0000 7fee11643700 10 mon.a@0(leader).auth v0 encode_pending v 1
```

but then before that commits,

```
2019-07-02T18:49:22.760+0000 7fee0de3c700 10 mon.a@0(leader).paxos(paxos refresh c 1..1) doing requested bootstrap
```

and then

```
2019-07-02T18:49:22.774+0000 7fee11643700 10 mon.a@0(leader).auth v0 create_initial -- creating initial map
2019-07-02T18:49:22.774+0000 7fee11643700 10 cephx keyserver: _check_rotating_secrets
2019-07-02T18:49:22.774+0000 7fee11643700 10 mon.a@0(leader).auth v0 check_rotate updated rotating
```

but that second time the `_check_rotating_secrets` doesn't add any new keys. probably it is not cleanly resetting its state to empty after the bootstrap?

#### Related issues:

Copied to RADOS - Backport #40730: nautilus: mon: auth mon isn't loading full...	<b>Resolved</b>
Copied to RADOS - Backport #40731: luminous: mon: auth mon isn't loading full...	<b>Rejected</b>
Copied to RADOS - Backport #40732: mimic: mon: auth mon isn't loading full Ke...	<b>Resolved</b>

#### History

##### #1 - 07/02/2019 11:05 PM - Sage Weil

- Status changed from *In Progress* to *Fix Under Review*
- Assignee deleted (*Sage Weil*)

<https://github.com/ceph/ceph/pull/28850>

##### #2 - 07/02/2019 11:05 PM - Sage Weil

- Backport set to *nautilus, mimic, luminous*

<https://github.com/ceph/ceph/pull/28850>

##### #3 - 07/09/2019 02:21 PM - Kefu Chai

- Status changed from *Fix Under Review* to *Pending Backport*

##### #4 - 07/11/2019 01:58 PM - Nathan Cutler

- Pull request ID set to *28850*

##### #5 - 07/11/2019 01:58 PM - Nathan Cutler

- Copied to Backport #40730: *nautilus: mon: auth mon isn't loading full KeyServerData after restart added*

##### #6 - 07/11/2019 01:58 PM - Nathan Cutler

- Copied to Backport #40731: *luminous: mon: auth mon isn't loading full KeyServerData after restart added*

##### #7 - 07/11/2019 01:58 PM - Nathan Cutler

- Copied to Backport #40732: *mimic: mon: auth mon isn't loading full KeyServerData after restart added*

##### #8 - 01/27/2021 07:09 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".