

fs - Bug #21848

client: re-expand admin_socket metavariables in child process

10/19/2017 07:51 AM - Zhi Zhang

Status:	Resolved	Start date:	10/19/2017
Priority:	High	Due date:	
Assignee:	Zhi Zhang	% Done:	0%
Category:	Administration/Usability	Estimated time:	0.00 hour
Target version:	v14.0.0	Affected Versions:	
Source:	Community (dev)	ceph-qa-suite:	
Tags:		Component(FS):	ceph-fuse
Backport:		Labels (FS):	
Regression:	No	Pull request ID:	
Severity:	3 - minor		
Reviewed:			

Description

The default value of admin_socket is \$run_dir/\$cluster-\$name.asok. If mounting multiple ceph-fuse instances on the same host with the same entity id, all can be mounted successfully but admin socket will fail as follows.

```
2017-10-19 14:59:00.631443 7ff6b3999040 -1 asok(0x7ff6bdda8380) AdminSocketConfigObs::init: failed
: AdminSocket::bind_and_listen: failed to bind the UNIX domain socket to '/var/run/ceph/ceph-clien
t.admin.asok': (17) File exists
```

Then we configure admin_socket to \$run_dir/\$cluster-\$name-\$pid.asok in ceph.conf file to avoid above failure. It works but the value of \$pid is not the final ceph-fuse process id. The value is the parent process id which will exit quickly. So it is still difficult to identify which asok file belongs to which ceph-fuse instance.

```
$ sudo ceph-fuse /mnt/ceph-fuse/
2017-10-19 15:00:08.005618 7efe6c2dd040 -1 init, newargv = 0x7efe76052760 newargc=9
ceph-fuse[89487]: starting ceph client
ceph-fuse[89487]: starting fuse
```

```
$ ls -l /var/run/ceph/ceph-client.admin-89485.asok
srwxr-xr-x 1 root root 0 Oct 19 15:00 /var/run/ceph/ceph-client.admin-89485.asok
```

The reason is that all the option's metavariables are expanded early, especially in ceph-fuse case, they are expanded in parent process however ceph-fuse is started in child process. So we try to re-expand admin_socket metavariables in child process if it is defined in conf file.

History

#1 - 10/19/2017 07:52 AM - Zhi Zhang

<https://github.com/ceph/ceph/pull/18393>

#2 - 10/23/2017 11:46 AM - Nathan Cutler

- Status changed from New to Need Review

#3 - 10/25/2017 08:35 PM - Patrick Donnelly

- Status changed from Need Review to In Progress

- Assignee set to Zhi Zhang

- Backport set to luminous

Zhi, please revisit this issue as the fix in <https://github.com/ceph/ceph/pull/18393> must be reverted due to the reasons trociny outlined.

#4 - 04/09/2018 09:05 PM - Patrick Donnelly

- Subject changed from *ceph-fuse: re-expand admin_socket metavariables in child process to client: re-expand admin_socket metavariables in child process*

- Category set to *Administration/Usability*

- Target version set to *v14.0.0*

- Backport changed from *luminous* to *luminous,mimic*

- Component(FS) *Client* added

- Component(FS) *deleted (ceph-fuse)*

#5 - 04/13/2018 09:48 AM - Zhi Zhang

Hi Patrick,

Sorry for missing this for a long time. I will take a look recently to see whether there is a better fixing way.

#6 - 04/20/2018 03:52 AM - Zhi Zhang

<https://github.com/ceph/ceph/pull/21544>

Patrick, could you pls take a look at this new fix? Now it is not only for admin_socket, but also for those options containing \$pid. Thanks.

#7 - 04/20/2018 04:33 PM - Patrick Donnelly

- Status changed from *In Progress* to *Need Review*

- Target version changed from *v14.0.0* to *v13.0.0*

- Backport changed from *luminous,mimic* to *luminous*

- Component(FS) *ceph-fuse* added

- Component(FS) *deleted (Client)*

#8 - 05/03/2018 07:37 PM - Patrick Donnelly

- Priority changed from *Normal* to *High*

- Target version changed from *v13.0.0* to *v14.0.0*

- Backport changed from *luminous* to *luminous,mimic*

#9 - 06/26/2018 09:08 PM - Patrick Donnelly

- Status changed from *Need Review* to *Resolved*

- Backport *deleted (luminous,mimic)*