

devops - Fix #16440

[initscripts]: systemd-run is not needed in initscripts

06/23/2016 08:36 AM - Vladislav Odintsov

Status: Resolved	% Done: 100%
Priority: Normal	
Assignee: Vladislav Odintsov	
Category:	
Target version:	
Source: Community (user)	Affected Versions:
Tags:	ceph-qa-suite:
Backport: jewel,hammer	Pull request ID:
Reviewed:	Crash signature:
Description	
systemd-run logic in initscripts was introduced because of ticket http://tracker.ceph.com/issues/7627 . If we have systemd-based distro, we should use systemd unit files from systemd directory to start/stop ceph daemons. Otherwise, daemon() from /etc/init.d/functions on systemd distro starts service in system.slice and everything works well for case, for example, when we use hammer on RH7. With this code it will start daemon with daemon() function from init.d/functions. systemd-run can not be run on non-systemd distros, so it's not needed on SysV systems. I've performed successful testing with hammer on RH6, RH7 systems with ceph-osd, ceph-mon, ceph-radosgw services.	
Related issues:	
Copied to devops - Backport #16441: jewel: [initscripts] systemd-run is not n...	Resolved
Copied to devops - Backport #16442: hammer: [initscripts]: systemd-run is not...	Resolved

History

#1 - 06/23/2016 08:39 AM - Vladislav Odintsov

<https://github.com/ceph/ceph/pull/9871>

backport tickets and PRs:

jewel:

<http://tracker.ceph.com/issues/16441>

<https://github.com/ceph/ceph/pull/9872>

hammer:

<http://tracker.ceph.com/issues/16442>

<https://github.com/ceph/ceph/pull/9873>

#2 - 06/23/2016 01:24 PM - Nathan Cutler

- Copied to Backport #16441: jewel: [initscripts] systemd-run is not needed in initscripts added

#3 - 06/23/2016 01:26 PM - Nathan Cutler

- Status changed from New to Fix Under Review

#4 - 06/23/2016 01:28 PM - Nathan Cutler

@Vladimir: If any changes are made to <https://github.com/ceph/ceph/pull/9871> these will have to be manually reflected in the backport PRs you have open. That's why we usually wait until the master PR has been merged before staging backports.

#5 - 06/23/2016 01:48 PM - Vladislav Odintsov

@Nathan, okay, no problem, I'll wait until merge and update backport PRs.

#6 - 07/11/2016 05:44 AM - Kefu Chai

- Status changed from Fix Under Review to Resolved

#7 - 07/11/2016 09:39 AM - Vladislav Odintsov

Need backports:

hammer

jewel

kraken

#8 - 07/11/2016 10:45 AM - Vladislav Odintsov

Backport to kraken: <http://tracker.ceph.com/issues/16652>

#9 - 07/11/2016 11:38 AM - Nathan Cutler

@Vladislav: Kraken backporting has not started yet. The "master" and "kraken" branches are currently in lock-step. See <http://docs.ceph.com/docs/master/dev/#what-is-merged-where-and-when> for general info on this topic.

#10 - 07/14/2016 10:48 AM - Nathan Cutler

- Copied to Backport #16442: hammer: [initscripts]: systemd-run is not needed in initscripts added

#11 - 07/28/2016 07:48 PM - Loic Dachary

- Status changed from Resolved to Pending Backport

#12 - 07/28/2016 08:40 PM - Vladislav Odintsov

More issue description.

My initial problem was connected with hammer radosgw on el7. I couldn't run radosgw under non-privileged user without shell (/sbin/nologin), because you can't ask su or sudo or runuser to start process from user without shell - these binaries do need shell. But daemon() from /etc/init.d/functions or start_daemon() from /lib/lsb/init-functions don't need it and are written for this special usage.

I decided to write 2 patches: for init-ceph and for init-radosgw to keep start of these daemons with initscripts in the same style.

The problem:

```
~]# getent passwd radosgw
radosgw:x:1003:1003:./var/lib/ceph/radosgw/./sbin/nologin
~]# /bin/su radosgw -c '/bin/radosgw -n client.radosgw.i-c5ef3fa1'
This account is currently not available.
~]# ps aux | grep rados
root      11402  0.0  0.0 112648   968 pts/0    S+   11:32   0:00 grep --color=auto rados
```

Right way (if we don't use systemd-units on el7 hammer):

```
~]# . /etc/init.d/functions
~]# daemon --user="radosgw" "/bin/radosgw -n client.radosgw.i-c5ef3fa1"
~]#
[ OK ]
~]# ps aux | grep rados
radosgw  11757  1.0  0.2 2151224 9468 ?        Ssl  11:34   0:00 /bin/radosgw -n client.radosgw.i-c5ef3fa1
root     11893  0.0  0.0 112648   968 pts/0    S+   11:34   0:00 grep --color=auto rados
```

Anyway, even if you don't use systemd-run, you can work through systemctl:

```

~]# systemctl status ceph-radosgw
● ceph-radosgw.service - LSB: radosgw RESTful rados gateway
   Loaded: loaded (/etc/rc.d/init.d/ceph-radosgw)
   Active: active (running) since Fri 2016-06-24 11:44:04 MSK; 3s ago
     Docs: man:systemd-sysv-generator(8)
  Process: 8023 ExecStop=/etc/rc.d/init.d/ceph-radosgw stop (code=exited, status=0/SUCCESS)
  Process: 20117 ExecReload=/etc/rc.d/init.d/ceph-radosgw reload (code=exited, status=0/SUCCESS)
  Process: 13781 ExecStart=/etc/rc.d/init.d/ceph-radosgw start (code=exited, status=0/SUCCESS)
   CGroup: /system.slice/ceph-radosgw.service
           └─13805 /bin/radosgw -n client.radosgw.i-c5ef3fa1

Jun 24 11:44:04 i-c5ef3fa1 systemd[1]: Starting LSB: radosgw RESTful rados gateway...
Jun 24 11:44:04 i-c5ef3fa1 ceph-radosgw[13781]: Starting client.radosgw.i-c5ef3fa1...
Jun 24 11:44:04 i-c5ef3fa1 runuser[13801]: pam_unix(runuser:session): session opened for user radosgw by (uid=0)
Jun 24 11:44:04 i-c5ef3fa1 systemd[1]: Started LSB: radosgw RESTful rados gateway.
Jun 24 11:44:04 i-c5ef3fa1 ceph-radosgw[13781]: [ OK ]

```

About backports:

It's very needed in EL7 hammer. There is no ability to start daemon under non-privileged user. Since el7 hammer doesn't provide packaged systemd units, and changing packaging is not a good idea in minor release, because it can break the automation and other aspects for users.

I think these changes should be in all supported branches (hammer, jewel) to keep these scripts in unified style - but this is already not a technical issue.

#13 - 08/25/2016 11:40 AM - Loic Dachary

- Copied to deleted (Backport #16441: jewel: [initscripts] systemd-run is not needed in initscripts)

#14 - 08/25/2016 11:42 AM - Loic Dachary

- Backport changed from hammer,jewel to hammer

See <http://tracker.ceph.com/issues/16441#note-9>

#15 - 11/23/2016 09:57 PM - Nathan Cutler

- Project changed from Ceph to devops

- Status changed from Pending Backport to Resolved

- Backport changed from hammer to jewel,hammer

#16 - 11/23/2016 09:57 PM - Nathan Cutler

- Copied to Backport #16441: jewel: [initscripts] systemd-run is not needed in initscripts added

#17 - 11/23/2016 09:57 PM - Nathan Cutler

- Copied to Backport #16442: hammer: [initscripts]: systemd-run is not needed in initscripts added

#18 - 11/23/2016 09:57 PM - Nathan Cutler

- % Done changed from 90 to 100