SELinux - Bug #19053

Proxy continues to write to deleted file after log rotation

03/28/2017 12:01 PM - Jason Nance

Status:	Closed		
Priority:	Normal		
Assignee:	Lukas Zapletal		
Category:	Smart proxy		
Target version:			
Difficulty:		**Fixed in Releases:**	1.23.0
Triaged:	No	**Found in Releases:**	
Bugzilla link:		**Red Hat JIRA:**	
Pull request:	https://github.com/theforeman/foreman-selinux/pull/75		

Description

After logrotate runs both foreman-proxy and squid continue to write to the rotated log, which is a deleted file handle.

```
$ sudo lsof /var/log | grep deleted
squid      1357          root    5u   REG  253,6     18393 2117707 /var/log/squid/cache.log-201703
27 (deleted)
squid      1359         squid    5u   REG  253,6     18393 2117707 /var/log/squid/cache.log-201703
27 (deleted)
ruby       1792 foreman-proxy    7w   REG  253,6     21154 1048669 /var/log/foreman-proxy/proxy.lo
g-20170325 (deleted)
```

This happens both on the Katello host and any Capsules.

I have tried manually running /bin/systemctl kill --signal=SIGUSR1 foreman-proxy >/dev/null 2>&1 || true (which is the postrotate script) but the file is not released and new file is not written to.

CentOS 7.3
Katello 3.2.2
Foreman 1.13.4
Squid 3.5.20
Logrotate 3.8.6

The only modification to /etc/logrotate.conf from the RPM-provided file is that compress has been uncommented.

Related issues:

Copied to SELinux - Bug #19223: AVC denied from logrotate sending signal to f...	**Needs design**	**04/07/2017**
Copied to foreman-tasks - Bug #19224: smart_proxy_dynflow_core.log not reopen...	**Closed**	**04/07/2017**

Associated revisions

Revision f9d17e7d - 06/26/2019 09:07 AM - Lukas Zapletal

Fixes #19053 - allow logrotate to send signals

History

#1 - 03/29/2017 02:09 PM - Justin Sherrill

- Project changed from Katello to Smart Proxy

- Category set to Packaging

- Priority changed from Urgent to High

The squid logrotate file appears to come from the squid rpm, so you may want to file a bug against centos/RHEL for that.

Moving this to smart-proxy packaging to handle the smart proxy logging.

#2 - 03/29/2017 02:14 PM - Jason Nance

Regarding the Squid LogRotate, this seems to be specific to the way Foreman is deploying it as I have non-Foreman Squid instances running on the same OS version in the same hosted environment which do not suffer from this issue.

#3 - 03/29/2017 03:48 PM - Anonymous

Could you check the system log to see if there any logrotate-related errors in it? Out of curiosity, is /var/lib filesystem mounted as RO? If so, please see https://bugzilla.redhat.com/show_bug.cgi?id=1272236.

#4 - 03/29/2017 03:48 PM - Anonymous

- Priority changed from High to Normal

#5 - 04/03/2017 04:07 PM - Jason Nance

I don't see any helpful logs. Here's what I can find:

```
Apr 03 03:06:01 sl1mmgplsat0001.ipa.gen.zone anacron[114112]: Job `cron.daily' started
Apr 03 03:06:02 sl1mmgplsat0001.ipa.gen.zone run-parts(/etc/cron.daily)[116790]: starting logrotate
Apr 03 03:06:05 sl1mmgplsat0001.ipa.gen.zone logrotate[116798]: ALERT exited abnormally with [1]
Apr 03 03:06:05 sl1mmgplsat0001.ipa.gen.zone run-parts(/etc/cron.daily)[116800]: finished logrotate
Apr 03 03:06:05 sl1mmgplsat0001.ipa.gen.zone run-parts(/etc/cron.daily)[116802]: starting man-db.cron
Apr 03 03:06:21 sl1mmgplsat0001.ipa.gen.zone run-parts(/etc/cron.daily)[116866]: finished man-db.cron
Apr 03 03:06:21 sl1mmgplsat0001.ipa.gen.zone anacron[114112]: Job `cron.daily' terminated (mailing output)
Apr 03 03:06:23 sl1mmgplsat0001.ipa.gen.zone anacron[114112]: Normal exit (1 job run)
```

/var/lib isn't a read-only filesystem (it is just a directory in /var, which is also not read-only):

```
/dev/mapper/vg_root-lv_var on /var type xfs (rw,relatime,seclabel,attr2,inode64,noquota)

# df -h
Filesystem                            Size  Used Avail Use% Mounted on
/dev/mapper/vg_root-lv_root           4.0G  2.4G  1.6G  60% /
devtmpfs                               12G     0   12G   0% /dev
tmpfs                                  12G   12K   12G   1% /dev/shm
tmpfs                                  12G  137M   12G   2% /run
tmpfs                                  12G     0   12G   0% /sys/fs/cgroup
/dev/sda1                             487M  156M  302M  35% /boot
/dev/mapper/vg_root-lv_var             75G   35G   41G  47% /var
/dev/mapper/vg_root-lv_opt           1014M  476M  539M  47% /opt
/dev/mapper/vg_root-lv_tmp           1014M  135M  880M  14% /tmp
/dev/mapper/vg_root-lv_home          1014M   34M  981M   4% /home
/dev/mapper/vg_root-lv_varlog         5.0G  2.4G  2.7G  48% /var/log
/dev/mapper/vg_root-lv_varlibmongodb   75G   34G   42G  45% /var/lib/mongodb
/dev/mapper/vg_root-lv_varlogaudit   1014M   64M  951M   7% /var/log/audit

drwxr-xr-x. 46 root root 4096 Mar 24 09:59 /var/lib
```

Running the postrotate in /etc/logrotate.d/foreman-proxy manually doesn't release the open file handler, either:

```
# lsof /var/log/ | grep delete
squid       1357           root    5u    REG  253,6     18393 2117707 /var/log/squid/cache.log-20170327 (deleted
)
squid       1359          squid    5u    REG  253,6     18393 2117707 /var/log/squid/cache.log-20170327 (deleted
)
ruby        1792 foreman-proxy    7w    REG  253,6     33954 1048669 /var/log/foreman-proxy/proxy.log-20170325 
(deleted)
# ps auxww | grep 1792
foreman+   1792  0.0  0.1 724000 43288 ?        Sl   Mar24   1:44 ruby /usr/share/foreman-proxy/bin/smart-prox
y
# systemctl kill --signal=SIGUSR1 foreman-proxy
# lsof /var/log/ | grep delete
squid       1357           root    5u    REG  253,6     18393 2117707 /var/log/squid/cache.log-20170327 (deleted
)
squid       1359          squid    5u    REG  253,6     18393 2117707 /var/log/squid/cache.log-20170327 (deleted
)
ruby        1792 foreman-proxy    7w    REG  253,6     33954 1048669 /var/log/foreman-proxy/proxy.log-20170325 
(deleted)
```

I noticed you change the priority on this report. Are you unable to reproduce this? 'Cause I would think that a bug which requires you to restart the application to release a deleted file handle would be high priority...

#6 - 04/04/2017 07:20 AM - Anonymous

Log rotation works as expected on my machine (fedora 25, logrotate 3.10.0).

Smart-proxy re-opens its log file on the first access to the log after it received a SIGUSR1. Try accessing smart-proxy (curl https://localhost:8443/version should do the trick) and check if it has switched to the new log file. If you are running logrotate manually, make sure you aren't using it in dry-run mode (i.e "-d" switch).

#7 - 04/04/2017 10:47 AM - Jason Nance

I'm not running logrotate manually, I was just running the postscript (sending the SIGUSR1) manually to take logrotate out of the picture for troubleshooting. Sending the SIGUSR1 isn't working for me.

I have also tried sending the SIGUSR1 using kill instead of systemctl and foreman-proxy isn't releasing the log file:

```
[root@sl1mmgplsat0001 ~]# lsof /var/log/ | grep delete
squid       1357          root    5u   REG  253,6     18393 2117707 /var/log/squid/cache.log-20170327 (deleted
)
squid       1359         squid    5u   REG  253,6     18393 2117707 /var/log/squid/cache.log-20170327 (deleted
)
ruby        1792 foreman-proxy    7w   REG  253,6     33954 1048669 /var/log/foreman-proxy/proxy.log-20170325
(deleted)
[root@sl1mmgplsat0001 ~]# kill -s SIGUSR1 1792
[root@sl1mmgplsat0001 ~]# lsof /var/log/ | grep delete
squid       1357          root    5u   REG  253,6     18393 2117707 /var/log/squid/cache.log-20170327 (deleted
)
squid       1359         squid    5u   REG  253,6     18393 2117707 /var/log/squid/cache.log-20170327 (deleted
)
ruby        1792 foreman-proxy    7w   REG  253,6     33954 1048669 /var/log/foreman-proxy/proxy.log-20170325
(deleted)
```

What other info can I provide to help track this down?

#8 - 04/04/2017 10:51 AM - Anonymous

> Sending the SIGUSR1 isn't working for me

Did you read the explanation of how smart-proxy handles log rotation? Did you try accessing it after sending SIGUSR1 to it?

#9 - 04/04/2017 11:18 AM - Jason Nance

Yes, I saw. My Smart Proxy is fairly actively used and as you can see from the output that file has been opened since around March 24. I'm 100% certain that the Smart Proxy has been accessed multiple times in the past 10 days.

Outside of the /version URL (which 404s for me, BTW), what else should trigger this (hosts checking in, content sync, content view changes)?

I tried browsing to Infrastructure -> Smart Proxies and then clicking around on the Smart Proxy tabs and that seemed to have worked for the Smart Proxy on my Katello server (where I just ran the kill), but doing the same for my Capsule (without manually running the kill) didn't work. Is it possible that there is a time limit on this or something?

#10 - 04/04/2017 11:31 AM - Anonymous

> Outside of the /version URL (which 404s for me, BTW), what else should trigger this

/version should always be accessible, via http, https, or both. If it's not available then smart-proxy isn't up. Any request to smart proxy should be sufficient, esp. if you have debug-level logging enabled (I'm not sure if listing proxies in Foreman is sufficient though). I'm not sure if any of content-related operations would suffice either, as they are handled by pulp with little involvement from proxy. To make sure, try getting a response from /version or /features urls. Depending on the configuration, you may need to use correct ip address/fqdn/protocol or combination of thereof.

> Is it possible that there is a time limit on this or something?

Not sure what you mean by this. When SIGUSR1 is caught, an internal flag signalling the need to reopen the log file is raised. It is reset when (and only then) the log file has been reopened.

#11 - 04/04/2017 12:06 PM - Jason Nance

Dmitri Dolguikh wrote:

> > Outside of the /version URL (which 404s for me, BTW), what else should trigger this
>
> /version should always be accessible, via http, https, or both. If it's not available then smart-proxy isn't up.

On my Katello host:

```
# wget --no-check-certificate https://localhost:8443/version
--2017-04-04 10:36:01--  https://localhost:8443/version
Resolving localhost (localhost)... ::1, 127.0.0.1
Connecting to localhost (localhost)|::1|:8443... connected.
    WARNING: certificate common name ‘sl1mmgplsat0001.ipa.gen.zone’ doesn't match requested host name ‘localho
st’.
HTTP request sent, awaiting response... 404 Not Found
2017-04-04 10:36:01 ERROR 404: Not Found.
```

On my Capsule:

```
# wget --no-check-certificate https://localhost:8443/version
--2017-04-04 10:38:01--  https://localhost:8443/version
Resolving localhost (localhost)... ::1, 127.0.0.1
Connecting to localhost (localhost)|::1|:8443... connected.
WARNING: no certificate subject alternative name matches
        requested host name ‘localhost’.
HTTP request sent, awaiting response... 404 Not Found
2017-04-04 10:38:01 ERROR 404: Not Found.
```

I also tried replacing "localhost" with the FQDN of the server and tried substituting "/version" with "/features". :-\

> > Is it possible that there is a time limit on this or something?
>
> Not sure what you mean by this. When SIGUSR1 is caught, an internal flag signalling the need to reopen the log file is raised. It is reset when (and only then) the log file has been reopened.

What I meant is if that re-open had to be triggered in a certain time period else the flag gets cleared. It seems that this flag either isn't getting set or isn't getting triggered.

#12 - 04/04/2017 12:51 PM - Anonymous

re: 404's -- check smart-proxy configuration (in /etc/foreman-proxy/settings.yml), it will have ip address(es) and ports on which proxy is listening.

> What I meant is if that re-open had to be triggered in a certain time period else the flag gets cleared. It seems that this flag either isn't getting set or isn't getting triggered.

No, there's no time limit or anything of the sort. The mechanics of catching the signal is trivial too (and the issue has never been reported before), so I'm very sceptical this is a bug. Based on description of how the capsule is being used, I don't think you are hitting smart-proxy itself (pulp is a separate process), and it's quite possible that the log file simply hasn't been reopened. Once you get a response from /version or /features urls, please check if the log file has been reopened.

#13 - 04/06/2017 05:43 PM - Jason Nance

Okay, the correct URL is:

https://localhost:9090/version

However, this does not trigger the rotation:

```
# wget -O- --no-check-certificate https://localhost:9090/version
--2017-04-06 16:42:23--  https://localhost:9090/version
Resolving localhost (localhost)... ::1, 127.0.0.1
Connecting to localhost (localhost)|::1|:9090... failed: Connection refused.
Connecting to localhost (localhost)|127.0.0.1|:9090... connected.
WARNING: no certificate subject alternative name matches
        requested host name ‘localhost’.
HTTP request sent, awaiting response... 200 OK
Length: 185 [application/json]
Saving to: ‘STDOUT’

 0% [                                       ] 0           --.-K/s              {"version":"1.13.4","modules":{
"pulp":"1.3.0","openscap":"0.6.0","dynflow":"0.1.5","ssh":"0.1.4","tftp":"1.13.4","puppetca":"1.13.4","puppet"
:"1.13.4","realm":"1100%[======================================>] 185         --.-K/s   in 0.04s

2017-04-06 16:42:23 (4.60 KB/s) - written to stdout [185/185]

# lsof /var/log/ | grep delete                          ruby        1792 foreman-proxy    7w   REG  253,6
12434 1050864 /var/log/foreman-proxy/proxy.log-20170405 (deleted)
```

#14 - 04/06/2017 05:46 PM - Jason Nance

Jason Nance wrote:

> However, this does not trigger the rotation:

I meant to say "the re-open". I understand that this isn't actually rotating the logs. The logrotate script is doing that.

#15 - 04/07/2017 05:17 AM - Anonymous

I'm not sure what's going on there. I still can't reproduce this, and I don't even see ruby vm holding log files open (lsof shows nothing). I would suggest asking a question about this in #theforeman irc channel, perhaps other users has seen this issue before?

#16 - 04/07/2017 08:50 AM - Jason Nance

I'll ask in IRC but I previously posted on the email list without a response.

Do you have an EL 7 box with Katello laying around to check this? Between production and lab I have 3 Katello hosts and 1 Capsule and all 4 instances are experiencing this. Of course, they are all on CentOS 7 and have the same plugins enabled. Is it possible that it is related to the Ruby version or a plugin?

#17 - 04/07/2017 09:05 AM - Anonymous

> Do you have an EL 7 box with Katello laying around to check this?

Not at the moment. I'm not sure what the issue there is; I couldn't replicate this with different versions of ruby (2.0.0 and 2.3.1). If you are seeing this with Katello too (which runs in Rails in a separate from smart-proxy VM), then it would suggest that the issue is outside of smart-proxy. You mentioned that you are seeing the same problem with squid, which would suggest that the issue isn't with ruby VM, but elsewhere (logrotate perhaps?).

#18 - 04/07/2017 09:23 AM - Anonymous

Try sending SIGUSR1 to smart-proxy via kill (as opposed to through systemctl)? See if that works any better?

#19 - 04/07/2017 09:59 AM - Jason Nance

It's better in a way...

```
[root@sl2mmgplsat0001 ~]# lsof /var/log | grep deleted
ruby      121892 foreman-proxy    7w   REG  253,5      1043 1049919 /var/log/foreman-proxy/smart_proxy_dynflow
_core.log-20170407 (deleted)
ruby      122001 foreman-proxy    7w   REG  253,5     13744 1049879 /var/log/foreman-proxy/proxy.log-20170407
(deleted)
```

```
[root@sl2mmgplsat0001 ~]# /bin/systemctl kill --signal=SIGUSR1 foreman-proxy
```

```
[root@sl2mmgplsat0001 ~]# lsof /var/log | grep deleted
ruby      121892 foreman-proxy    7w   REG  253,5      1043 1049919 /var/log/foreman-proxy/smart_proxy_dynflow
_core.log-20170407 (deleted)
```

The old proxy.log was released, but not the smart_proxy_dynflow_core.log (which I don't recall seeing before). Is there a different way to trigger the reopen of smart_proxy_dynflow_core.log?

I did, however, finally find this in the audit.log from early this morning:

```
type=USER_AVC msg=audit(1491468305.905:16288): pid=1 uid=0 auid=4294967295 ses=4294967295 subj=system_u:system
_r:init_t:s0 msg='avc:  denied  { stop } for auid=0 uid=0 gid=0 path="/usr/lib/systemd/system/foreman-proxy.se
rvice" cmdline="/bin/systemctl kill --signal=SIGUSR1 foreman-proxy" scontext=system_u:system_r:logrotate_t:s0-
s0:c0.c1023 tcontext=system_u:object_r:systemd_unit_file_t:s0 tclass=service  exe="/usr/lib/systemd/systemd" s
auid=0 hostname=? addr=? terminal=?'
```

It looks like SELinux is preventing logrotate from sending the kill. For some reason that wasn't in the previous audit.log (or I missed it).

So assuming there is a different way to trigger the reopen of smart_proxy_dynflow_core.log this looks to be an SELinux issue. If there isn't a different way to trigger the reopen of smart_proxy_dynflow_core.log it looks like I may have two separate issues.

#20 - 04/07/2017 10:44 AM - Anonymous

I took a quick look at dynflow_core, but not sure if/how they handle log rotation. The easiest thing would be to ask this question in #theforeman.

#21 - 04/07/2017 10:46 AM - Anonymous

- Tracker changed from Bug to Support

- Category changed from Packaging to Core

I'm going to change this to support; Please file file separate bug reports with Foreman SELinux and/or Smart-Proxy Dynflow module (if you think the issues above are bugs).

#22 - 04/07/2017 11:10 AM - Jason Nance

- Copied to Bug #19223: AVC denied from logrotate sending signal to foreman-proxy unit added

#23 - 04/07/2017 11:24 AM - Jason Nance

- Copied to Bug #19224: smart_proxy_dynflow_core.log not reopened during logrotate added

#24 - 04/27/2017 10:48 AM - Anonymous

- Status changed from New to Resolved

#25 - 09/19/2017 07:56 AM - Lukas Zapletal

- Tracker changed from Support to Bug

- Status changed from Resolved to Assigned

- Assignee set to Lukas Zapletal

So the problem here is the standard practice is to send SIGHUP to daemons to release log files after rotating, base RHEL SELinux policy includes rules to allow logrotate to send this via systemctl reload command. Our proxy only supports SIGUSR1 which is not standard and therefore the only way supporting this is by doing changes in our SELinux policy. This includes creating two types, domain, macro, file contexts - this is too complicated. It is much better to follow what is expected and change our service to respond to reload properly, then we can change our logrotate script as well.

Therefore I am opening this and making the change in proxy to properly reopen logs when "reload" init command is sent.

#26 - 09/19/2017 07:59 AM - The Foreman Bot

- Status changed from Assigned to Ready For Testing

- Pull request https://github.com/theforeman/smart-proxy/pull/544 added

#27 - 09/19/2017 08:05 AM - Lukas Zapletal

- Pull request https://github.com/theforeman/foreman-packaging/pull/1814 added

Two PRs are associated with this one.

#28 - 09/26/2017 06:46 AM - Lukas Zapletal

- Project changed from Smart Proxy to SELinux

- Category deleted (Core)

- Status changed from Ready For Testing to New

- Pull request deleted (https://github.com/theforeman/foreman-packaging/pull/1814, https://github.com/theforeman/smart-proxy/pull/544)

Unfortunately, proxy folks insist on using USR1 signal, therefore the only way is to modify our SELinux policy and introduce new domain.

https://github.com/theforeman/smart-proxy/pull/544
https://github.com/theforeman/foreman-packaging/pull/1814

#29 - 12/19/2017 09:54 AM - Lukas Zapletal

- Project changed from SELinux to Smart Proxy

- Category set to Packaging

I misread the PR, folks are not against SIGHUP, they did not like "reload" systemd action. Moving to smart-proxy and creating different proposal.

#30 - 12/19/2017 09:55 AM - The Foreman Bot

- Status changed from New to Ready For Testing

- Pull request https://github.com/theforeman/smart-proxy/pull/556 added

#31 - 12/19/2017 12:32 PM - Lukas Zapletal

- Project changed from Smart Proxy to SELinux

- Category deleted (Packaging)

- Status changed from Ready For Testing to New

- Pull request deleted (https://github.com/theforeman/smart-proxy/pull/556)

Ok, sorry for the confusion. I've confirmed with SELinux RH team that this is bug in RHEL policy and it will be fixed, I filed a bug: https://bugzilla.redhat.com/show_bug.cgi?id=1527522

@Jason can you please try this workaround, compile a custom policy with this rule:

allow logrotate_t systemd_unit_file_t:service stop;

That should fix it. I am unable to reproduce on all my RHEL 7.4 systems, but I will prepare a PR that will add this rule into our policy until this is fixed in RHEL 7.6+.

#32 - 12/19/2017 12:38 PM - The Foreman Bot

- Status changed from New to Ready For Testing

- Pull request https://github.com/theforeman/foreman-selinux/pull/75 added

#33 - 12/27/2017 09:30 AM - Radosław Piliszek

@Lukas:

I believe this is a too broad change. It would be better to introduce a new type for foreman-proxy service.

7.4 **is** affected.

#34 - 01/01/2018 10:31 AM - Lukas Zapletal

> I believe this is a too broad change. It would be better to introduce a new type for foreman-proxy service.

RH SELinux team confirmed me they are adding this rule into base policy in RHEL 7.5 so you will eventually get it. It must be stop because that's how SELinux hooks are implemented in systemd, a bit clunky, yeah.

#35 - 01/01/2018 10:47 AM - Radosław Piliszek

I see. Thanks for clarifying that. My hope is just that it will be implemented in such a way that it will not allow evil logrotate script to stop e.g. firewalld which would open all ports.

#36 - 01/02/2018 10:09 AM - Lukas Zapletal

Agree, this should be properly solved on the systemd/selinux level and not in all individual applications :-)

#37 - 01/04/2018 07:40 AM - Radosław Piliszek

Not exactly. My point is that it would be better if each application decided for itself whether it wants (needs) to be stopped/killed by logrotate via systemd.

Anyway:

There is already a rule:

```
allow logrotate_t domain : process signal ;
```

which essentially allows logrotate to send signals other than SIGKILL, SIGSTOP or SIGCHLD to virtually any other process (since all domain types should have domain attribute). So, well, adding systemd stopping would not be that bad (in fact only adding SIGKILL, SIGSTOP and SIGCHLD signals for systemd-managed services). :-)

#38 - 06/26/2019 09:07 AM - Tomer Brisker

- Fixed in Releases 1.23.0 added

#39 - 06/26/2019 10:01 AM - Anonymous

- Status changed from Ready For Testing to Closed

Applied in changeset f9d17e7d6f1ac1066b924f89490da12a400a4860.

#40 - 07/24/2019 03:00 PM - Amit Upadhye

- Category set to Smart proxy