Arvados - Bug #7228

[Crunch] crunch-dispatch should not create tmp files that break API server

09/07/2015 05:54 PM - Tom Clegg

Status:	Resolved	**Start date:**	09/07/2015
Priority:	Normal	**Due date:**	
Assigned To:	Peter Amstutz	**% Done:**	100%
Category:	Crunch	**Estimated time:**	0.00 hour
Target version:	2015-09-30 sprint		

Description

Problem

We have seen crunch-dispatch break API server as follows:

- Run as root, as described in docs
- Call some part of the API server's code base that uses Rails.cache
- Create files and directories in {Rails.root}/tmp/cache with owner=root and permissions that prohibit www-data from writing

After this has happened, API server (running as www-data) crashes when trying to update cached data.

This isn't very common because API server usually creates/updates a given cache item before crunch-dispatch does. But when it does happen, it's bad: for example, new groups can't be created because the group cache can't be updated.

The condition can be fixed temporarily by running arvados-api-server-upgrade.sh (it does chown -R on the tmp dir, among other things). However, this doesn't prevent it from happening again.

The real solution is #5162: refactor crunch-dispatch as an API client so it can't touch the API server Rails project at all.

Ideas

In the meantime, there might be an effective workaround, like running crunch-dispatch with umask=002 and the same GID as the API server process.

(Running crunch-dispatch with the same *UID* as the API server process would fix the cache permission issue, but at the cost of introducing other problems: crunch-dispatch needs to use sudo, and giving www-data passwordless sudo undermines the security benefit of running the web service as non-root in the first place.)

Immediate fix

- arvados-api-server-upgrade.sh already makes $WWW_OWNER the owner of tmp/ recursively. Extend it to chmod tmp/cache/ 2775.
- Extend crunch-dispatch to run with a 002 umask. The only other file it opens is its own lockfile, and it sets a specific 0644 mode for that, so this should only affect Rails cache files.
- Test using the procedure in note-5.
- Make sure the arvados-dev branch gets merged before the arvados branch, so we build a new package that includes both the new upgrade script and the new crunch-dispatch.

Subtasks:

Task # 7345: Review arvados branch 7228-crunch-dispatch-umask	**Resolved**
Task # 7374: Review 7228-group-writable-tmp-cache	**Resolved**

Related issues:

Related to Arvados - Bug #5162: [Crunch] crunch-dispatch should use the API i...	**Closed**

Associated revisions

Revision 469c824c - 09/23/2015 05:02 PM - Peter Amstutz

Merge branch '7228-group-writable-tmp-cache' refs #7228

Revision 469c824c - 09/23/2015 05:02 PM - Peter Amstutz

Merge branch '7228-group-writable-tmp-cache' refs #7228

Revision acd1241c - 09/24/2015 08:36 PM - Peter Amstutz

Merge branch '7228-crunch-dispatch-umask' closes #7228

History

#1 - 09/08/2015 09:38 AM - Brett Smith

- Category set to Crunch

Tom Clegg wrote:

> In the meantime, there might be an effective workaround, like running crunch-dispatch with umask=002 and the same GID as the API server process.

This seems to have some promise. Cache files currently in production have mode 0644, where www-data's umask is 022. It looks like the Rails cache code is not imposing a mode stricter than the umask.

> (Running crunch-dispatch with the same *UID* as the API server process would fix the cache permission issue, but at the cost of introducing other problems: crunch-dispatch needs to use sudo, and giving www-data passwordless sudo undermines the security benefit of running the web service as non-root in the first place.)

What if we took the other approach, and ran both as the crunch user? That's easy to rig up in Nginx: chown crunch: /var/www/arvados-api/current/config.ru.

#2 - 09/08/2015 02:06 PM - Brett Smith

- Target version set to Arvados Future Sprints

#3 - 09/15/2015 03:12 PM - Brett Smith

Brett Smith wrote:

> Tom Clegg wrote:
>
> > In the meantime, there might be an effective workaround, like running crunch-dispatch with umask=002 and the same GID as the API server process.
>
> This seems to have some promise. Cache files currently in production have mode 0644, where www-data's umask is 022. It looks like the Rails cache code is not imposing a mode stricter than the umask.

I dug into this to verify it. Here's Rails' code to write a cache file. Here's the atomic_write method called there. atomic_write specifically uses the default ownership and permissions of files in that directory. So setting tmp/cache setgid and umask 002 should be sufficient to fix the issue.

For what it's worth: the tracebacks we get from this refer to groups_for_user files. These are caches specifically created in the User model, calling methods on Rails cache. So yet another option might be to set up a different cache (another FileStore backed by a different directory?) when we're running crunch-dispatch, although I'm less sure how to do that nicely.

#4 - 09/15/2015 03:57 PM - Brett Smith

- Target version changed from Arvados Future Sprints to 2015-09-30 sprint

#5 - 09/15/2015 06:20 PM - Brett Smith

Manual test procedure:

- Start an API server without crunch-dispatch. Put at least one job in its queue.
- Stop the API server.
- Empty Rails' tmp/cache directory.
- Start crunch-dispatch.
- Confirm that cache files created by crunch-dispatch are writable by the Web server user (www-data on Debian).

#6 - 09/15/2015 06:20 PM - Brett Smith

- Story points set to 0.5

#7 - 09/15/2015 06:42 PM - Brett Smith

- Assigned To set to Brett Smith

#8 - 09/15/2015 07:36 PM - Brett Smith

- Assigned To changed from Brett Smith to Peter Amstutz

#9 - 09/15/2015 07:46 PM - Brett Smith

- Description updated

#10 - 09/23/2015 02:26 PM - Nico César

on arvados dev reviewing b3badf7..e48701c only 1 line change in

jenkins/arvados-api-server-extras/arvados-api-server-upgrade.sh

LGTM!

#11 - 09/23/2015 05:02 PM - Peter Amstutz

Test process:

1. Temporarily add "User.first.group_permissions" to #run method in crunch-dispatch to cause a cache to be written (I couldn't figure out the conditions that cause crunch-dispatch to create cache files consistently on startup)
2. Run sudo -u root -g peter crunch-dispatch.rb
3. This causes a groups_for_user file to be created.
4. Without umask:

   ```
   -rw-r--r-- 1 root peter  530 Sep 23 12:40 groups_for_user_bseta-tpzed-000000000000000
   ```

   API server breaks because it tries to delete groups_for_user_bseta-tpzed-000000000000000 and it can't
5. With umask:

   ```
   -rw-rw-r-- 1 root peter  531 Sep 23 12:41 groups_for_user_bseta-tpzed-000000000000000
   ```

   API server works because it is able to delete groups_for_user_bseta-tpzed-000000000000000

#12 - 09/23/2015 05:03 PM - Peter Amstutz

Branch is 7228-crunch-dispatch-umask

#13 - 09/24/2015 08:24 PM - Nico César

reviewed 27cb821..4a1fdbc

LGTM

#14 - 09/24/2015 08:40 PM - Peter Amstutz

- Status changed from New to Resolved

Applied in changeset arvados|commit:acd1241cec9260d54c2dca55785e309644334c41.