

RedHatInsights / insights-core Public[Code](#) [Issues](#) 131 [Pull requests](#) 10 [Actions](#) [Projects](#) [Wiki](#) [Security](#)

feat: improve temp directories #3878

Merged [xiangce](#) merged 1 commit into [RedHatInsights:master](#) from [ahitacat:temp-directory](#) on Aug 23, 2023

[Conversation](#) 11 [Commits](#) 1 [Checks](#) 0 [Files changed](#) 4



ahitacat commented [on Aug 9, 2023](#) • edited

Contributor

Check all that apply:

- Have you followed the guidelines in our Contributing document, including the instructions about commit messages?
- Is this PR to correct an issue?
- Is this PR an enhancement?

Complete Description of Additions/Changes:

This commit improve the use of temporary directories:

- The temporary directory will have a random name: `/var/tmp/insights-client-<random characters>`.
- It checks that the directory is not being in use before using it. If it existed, it will clean that directory and create a new one.
- Remove unused temporary directory and the function that used it: `archive_tmp_directory`.
- Change the cleanup mechanisms to be align with this change.
- Change the tests to be align with this change.
- Refactor the constants used.



ahitacat added the [client](#) label [on Aug 9, 2023](#)



ahitacat marked this pull request as ready for review [3 years ago](#)

ahitacat commented [on Aug 10, 2023](#)


Contributor

Author

[@ptoscano](#) Can you take a look at this?

 **ptoscano** reviewed [on Aug 17, 2023](#)

[View reviewed changes](#)

 **ptoscano** left a comment

Contributor

Generally LGTM (will need [@xiangce](#) review as well); I left one note about potential conflicts between instances.

`insights/client/archive.py` Outdated

Comment on lines 246 to 250

 **ptoscano** [on Aug 17, 2023](#)

Contributor

As I understand, there may be more than one `insights-client` run at the same time; in that case, because of this code there may be this situation:

- `insights-client` 1 starts: it creates its own `/var/tmp/insights-client-F00` temporary directory
- `insights-client` 2 starts as well: `cleanup_previous_archive()` cleans up all the existing `/var/tmp/insights-client*` stuff including `/var/tmp/insights-client-F00`
- `insights-client` 1 dies in sadness

What I'd do is removing only the bits older than 24h: this way even slow `insights-client` run will not suddenly get their own temporary directories removed under their feet, yet still cleaning old leftovers

 **xiangce** [on Aug 21, 2023](#)

Contributor

How about adding a restriction that only one instance of "insights-client" is allowed to run at a time?

 **ahitacat** [on Aug 21, 2023](#)

Contributor

Author

I am not sure if that is a good idea. If we added so, it will limit the following situations:

1. `insights-client` is running the scheduled collection and the administrator wants to check `insights-client` status.
2. The administrator want to run `collection` and also `malware` .
- 3.

I think that a better option is to allow this kind of situation, that was the first idea (before change the things in the temporary files)

ahitacat commented [on Aug 21, 2023](#)

Contributor

Author

Last changes pushed:

- The cleanup mechanism checks if the directories found with the determined prefix are older than 24 hours. Those that are older are removed.
- New tests.



 **ptoscano** reviewed [on Aug 21, 2023](#)

[View reviewed changes](#)



ptoscano left a comment

Contributor

There is at least one change needed (the variable used for `getmtime()`), and it should be OK.

Another possibility to make it a bit more Pythonic (IMHO) could be to use the list comprehension approach more, like the following untested code:

```
# list our temporary directories
directories = [
    os.path.join(constants.insights_tmp_path, directory)
    for directory in os.listdir(constants.insights_tmp_path)
    if directory.startswith(constants.insights_tmp_prefix)
]
# filter only directories modified last more than 24 hours ago
directories = [
    directory
    for directory in directories
    if os.path.isdir(directory) and os.path.getmtime(directory) <= one_day_ago
]
for directory in directories:
    logger.debug("Deleting previous archive %s", directory)
    shutil.rmtree(directory, True)
```





insights/client/archive.py Outdated



ptoscano on Aug 21, 2023

Contributor

this `directory` should most likely be `directory_path`

insights/client/archive.py Outdated



ptoscano on Aug 21, 2023

Contributor

if `directory` here is replaced by `os.path.join(constants.insights_tmp_path, directory)` , this avoids the need to do the `join()` in the loop



[feat: improve temp directories](#) ...

[ffe0aba](#)

ptoscano commented on Aug 22, 2023

Contributor

Thanks for the updates! IMHO LGTM now, I'll leave it to [@xiangce](#) for the final review.



xiangce approved these changes on Aug 23, 2023

[View reviewed changes](#)



xiangce left a comment

Contributor

It looks good to me too. And I tested the py26 on my local, it's PASSED.



xiangce merged commit [2286fbe](#) into [RedHatInsights:master](#) on Aug 23, 2023



xiangce pushed a commit that referenced this pull request on Aug 24, 2023



[feat: improve temp directories \(#3878\)](#) ...

[f2c6e11](#)



xiangce pushed a commit that referenced this pull request [on Sep 6, 2024](#)



`feat: improve temp directories (#3878)` ...

[7a05218](#)



xiangce pushed a commit that referenced this pull request [on Dec 11, 2025](#)



`feat: improve temp directories (#3878)` ...

[76ffb50](#)

[Sign up for free](#)

to join this conversation on GitHub. Already have an account? [Sign in to](#)

[comment](#)

Reviewers



xiangce



+1 more reviewer



ptoscano



Assignees

No one assigned

Labels

client

Projects

None yet

Milestone

No milestone

Development

Successfully merging this pull request may close these issues.

None yet

3 participants

