



链滴

ROTATING GC LOG FILES

作者: [fengjx](#)

原文链接: <https://ld246.com/article/1502711316726>

来源网站: [链滴](#)

许可协议: [署名-相同方式共享 4.0 国际 \(CC BY-SA 4.0\)](#)

Garbage Collection logs are essential artifacts to optimize application's performance and to avoid shoot complex memory problems. Garbage Collection logs can be generated in a particular file path by passing the "-Xloggc" JVM argument.

Example: `-Xloggc:/home/GCEASY/gc.log`

But the challenge to this approach is: whenever the application is restarted, old GC log file will be over-ridden by the new GC log file as the file path is same (i.e. /home/GCEASY/gc.log).

Thus you wouldn't be able to analyze the old GC logs that existed before restarting the application. Especially if the application has crashed or had certain performance problems then, you need old GC Logs for analysis.

Because of the heat of the production problem, most of the time, IT/DevOps team forgets to back up the old GC log file; A classic problem that happens again & again, that we all are familiar with (-). Most human errors can be mitigated through automation and this problem is no exception to it.

A simple strategy to mitigate this challenge is to write new GC log contents in a different file location. In this article 2 different strategies to do that are shared with you:

1. Suffix timestamp to GC Log file

If you can suffix the GC log file with the time stamp at which the JVM was restarted then, GC log file locations will become unique. Then new GC logs will not over-ride the old GC logs. It can be achieved as shown below:

UNIX:

```
-XX:+PrintGCDetails -XX:+PrintGCDateStamps -Xloggc:/home/GCEASY/gc-$(date +%Y_%m_%d-%H_%M).log
```

WINDOWS:

```
# Generating time stamp & replacing " with 0 and / with -
set file_suffix=%DATE:~-4%-DATE:~4,2%-DATE:~7,2%_TIME:~0,2%-TIME:~3,2%-TIME:~6,2%
set file_suffix=%file_suffix: =0%
set file_suffix=%file_suffix:/=-%
```

```
set JAVA_OPTS=%JAVA_OPTS% -XX:+PrintGCDetails -XX:+PrintGCDateStamps -Xloggc:c:\workspace\newUploads\gc-%file_suffix%.log
```

This strategy has one minor drawback:

a. Growing file size

Suppose if you don't restart your JVMs, then GC log file size can be growing to huge size. Because in this strategy new GC log files are created only when you restart the JVM. But this is not a major concern in my opinion, because one GC event only occupies few bytes. So typically GC log file size will not grow beyond a manageable point.

2. Use -XX:+UseGCLogFileRotation

Another approach is to use the JVM system properties:

```
-XX:+PrintGCDetails -XX:+PrintGCDateStamps -Xloggc:/home/GCEASY/gc.log -XX:+UseGCLogFileRotation -XX:NumberOfGCLogFiles=5 -XX:GCLogFileSize=2M
```

When `'-XX:-UseGCLogFileRotation'` is passed GC log rotation is enabled by the JVM itself.

`'-XX:NumberOfGCLogFiles'` sets the number of files to use when rotating logs, must be ≥ 1 . The rotated log files will use the following naming scheme, `.0`, `.1`, ..., `.n-1`.

`'-XX:GCLogFileSize'` defines the size of the log file at which point the log will be rotated, must be $\geq 8K$

But this strategy has few challenges:

a. Losing old GC Logs

Suppose if you had configured `-XX:NumberOfGCLogFiles=5` then, over a period of time, 5 GC log files will be created:

gc.log.0 — oldest GC Log content

gc.log.1

gc.log.2

gc.log.3

gc.log.4 — latest GC Log content

Most recent GC log contents will be written to `'gc.log.4'` and old GC log content will be present in `'gc.log.0'`.

When the application starts to generate more GC logs than the configured `'-XX:NumberOfGCLogFiles'` in this case 5, then old GC log contents in `gc.log.0` will be deleted. New GC events will be written to `gc.log.0`. It means you will end up not having all the generated GC logs. You will lose the full visibility of the events.

b. Mixed-up GC Logs

Suppose application has created 5 gc log files i.e.

gc.log.0

gc.log.1

gc.log.2

gc.log.3

gc.log.4

then, let's say you are restarting the application. Now new GC logs will be written to `gc.log.0` file and old GC log content will be present in `gc.log.1`, `gc.log.2`, `gc.log.3`, `gc.log.4` i.e.

gc.log.0 — GC log file content after restart

gc.log.1 — GC log file content before restart

gc.log.2 — GC log file content before restart

gc.log.3 — GC log file content before restart

gc.log.4 — GC log file content before restart

So your new GC log contents get mixed up with old GC logs. Thus to mitigate this problem you might have to move all the old GC logs to a different folder before you restart the application.

c. Forwarding GC logs to central location

In this approach, current active file to which GC logs are written is marked with extension “.current”. Example, if GC events are currently written to file ‘gc.log.3’ it would be named as: gc.log.3.current’.

If you want to forward GC logs from each server to a central location, then most DevOps engineers use ‘rsyslog’. However this file naming convention poses significant challenge to use ‘rsyslog’, as described in this blog.

d. Tooling

Now to analyze the GC log file using the GC tools such as (gceasy.io, GCViewer...), you will have to upload multiple GC log files instead of just one single GC Log file.

Conclusion

You can debate on which approach you want to take for rotating GC log files, but don’t debate on whether to rotate the GC log files or not. It will come very handy when need comes. You never know when need comes.

原文连接: <https://blog.gceasy.io/2016/11/15/rotating-gc-log-files/>