使用fluentbit收集日志，通过firehose打入S3

参考文档https://aws.amazon.com/cn/blogs/china/easy-aws-fargate-container-log-processing-with-aws-firelens/

https://github.com/aws-samples/amazon-ecs-firelens-examples

包含如下步骤

1. 创建app docker -- 使用firelens输出日志
2. 创建fluentbit docker -- 收集app docker中的指定目录下的日志，输出至firehose
3. 创建ecs task definition
4. 创建firehose datastream，打入S3

1. 创建app docker

使用如下docker配置，上传ecr构建app镜像

Dockerfile:

```
FROM alpine
COPY testlog.sh /bin/
RUN chmod +x /bin/testlog.sh
ENTRYPOINT ["/bin/testlog.sh"]
```

testlog.sh

```
#!/bin/sh

while :
do
   echo "{\"server_date\":\"2020-01-19\",\"hostname\":\"ip-172-31-43-24.cn-northwest-1.compute.internal\",\"pid\":5404,\"method\":\"POST\",\"clientIP\":\"10.11.12.13\",\"countryCode\":\"ID\",\"url\":\"/v1/mail/list\",\"status\":\"200\",\"latency\":7,\"length\":24,\"userId\":9536605,\"code\":20001}" >> /data/access.log
   echo "{\"server_date\":\"2020-01-19\",\"hostname\":\"ip-172-31-43-24.cn-northwest-1.compute.internal\",\"pid\":1000,\"method\":\"GET\",\"clientIP\":\"20.21.22.23\",\"countryCode\":\"ID\",\"url\":\"/v1/mail/list\",\"status\":\"500\",\"latency\":10,\"length\":12,\"userId\":1010001,\"code\":10001}" >> /data/error.log
sleep 10
done
```

2. 创建fluentbit docker

使用如下docker配置，上传ecr构建fluentbit镜像。因为需要使用自定义的fluentbit config文件，所以需要自定义fluentbit镜像，其中包含本地配置文件extra.conf。后面需要告诉fluentbit镜像使用该本地file作为配置文件，用于修改fluentbit使用的缺省配置文件（/fluent-bit/etc/fluent-bit.conf）

Dockerfile:

```
FROM amazon/aws-for-fluent-bit:latest
COPY extra.conf /extra.conf
```

Extra.conf

```
[SERVICE]
    Parsers_File /fluent-bit/parsers/parsers.conf
    Flush 1
    Grace 30

[INPUT]
    Name tail
    Path /data/access.log
    Tag access

[FILTER]
    Name parser
    Match *
    Key_Name log
    Parser json
    Reserve_Data True

[OUTPUT]
    Name firehose
    Match access
    region cn-northwest-1
    delivery_stream fluentbit-access
```

Input 为使用tail的方式收集/data/access.log内容，输出到名为fluentbit-access的firehose datasream中

3. 创建ecs task definition

- 根据目前使用模式，选择fargate

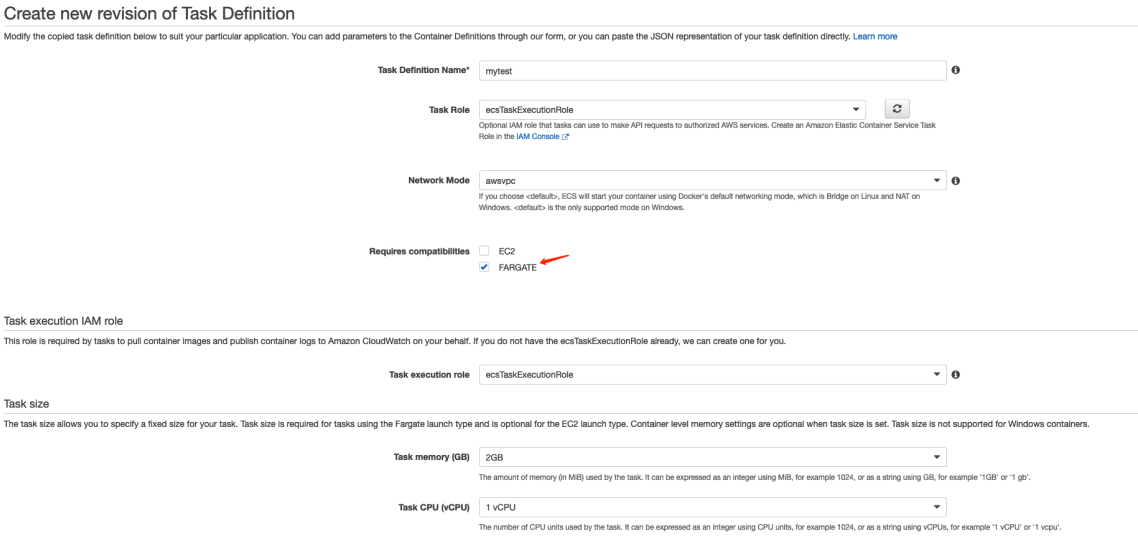

- 选择enable firelens integration，image使用上面创建的fluentbit 镜像

  创建名字为data的volume，用于两个container共享，app container将log输出到该volume，fluentbit从该volume读取日志发至firehose

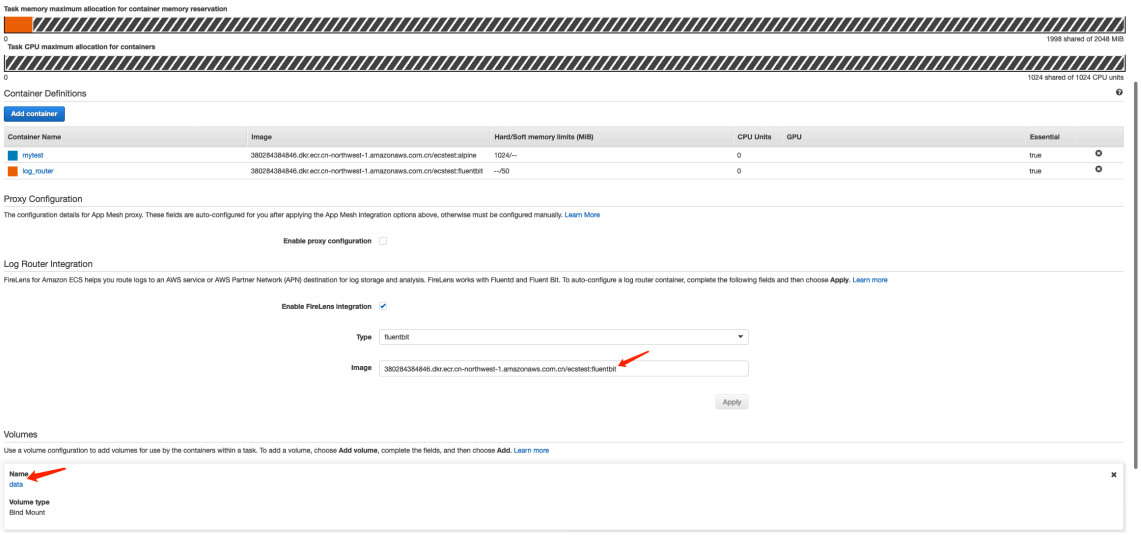

- 使用json格式，告知fluentbit使用extra.conf作为配置文件：

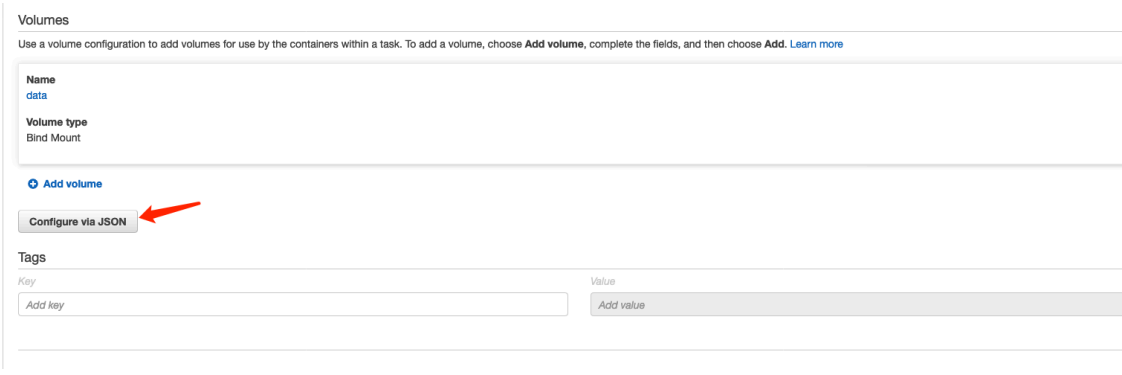

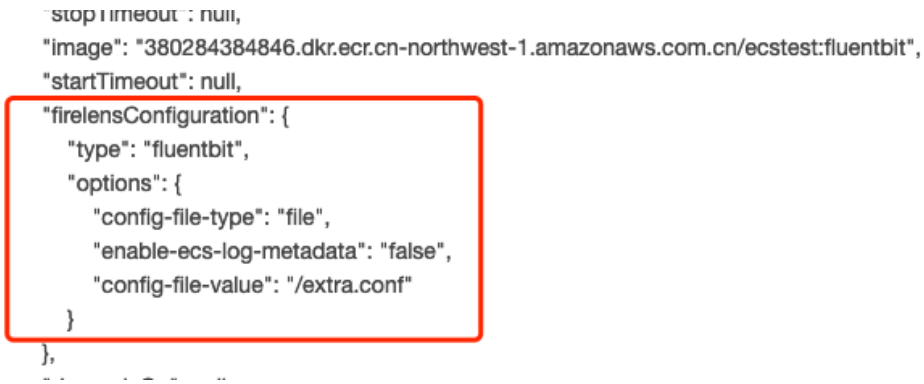


```
"firelensConfiguration": {
                "type": "fluentbit",
                "options": {
                    "config-file-type": "file",  ##可以为s3或file，对fargate只能用file
                    "enable-ecs-log-metadata": "false",
                    "config-file-value": "/extra.conf"
                }
            },
```

- 分别配置app container和fluentbit container的日志输出和文件挂载

  app container使用firelens的配置进行输出

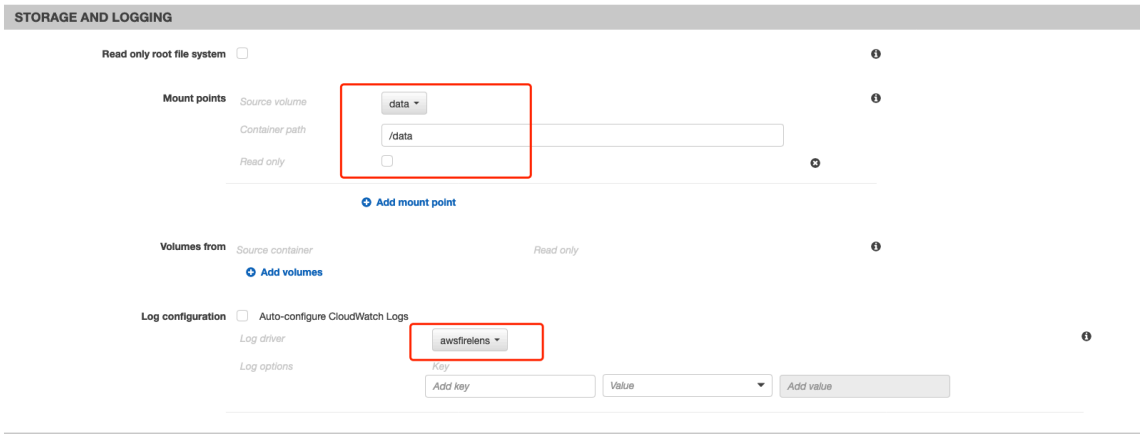

---

fluentbit container使用cloudwatch log进行日志输出，即fluentbit的运行情况从cloudwatch log进行观察：

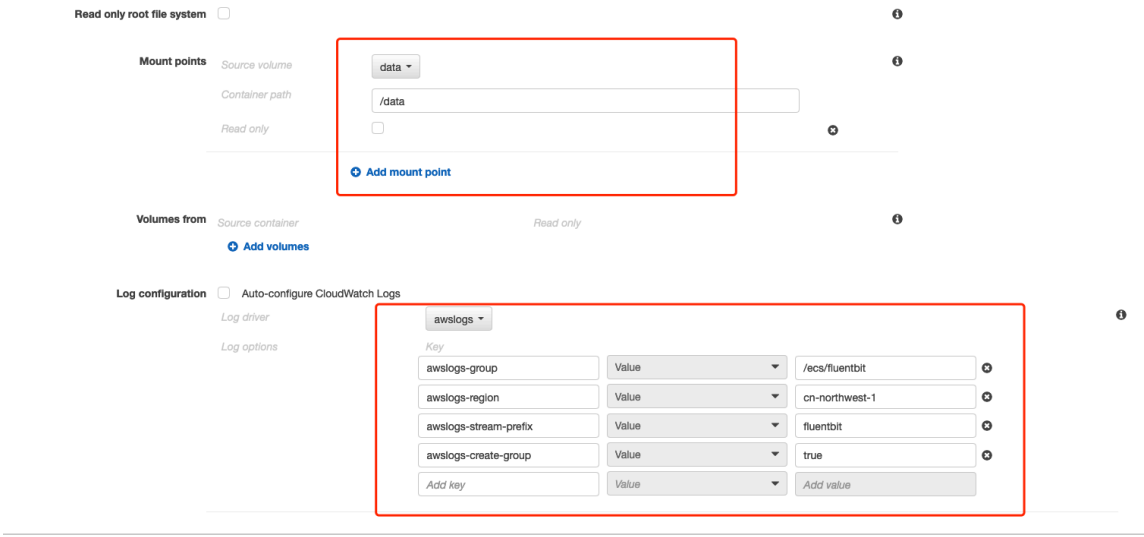

4. 创建firehose datastream

- 创建firehose datastream

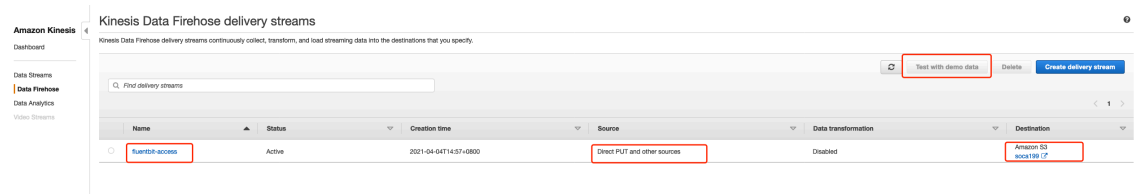

注意delivery stream应该使用fluentbit output中的delivery_stream相同的名称。

为防止权限有关的问题，使用自带的测试数据功能，先保证firehose可以将数据输出到s3

- 创建firehose endpoint

对firehose创建vpc endpoint，保证vpc内启动的fluentbit container可以访问到firehose服务。否则fluentbit container访问firehose时可能报访问firehose url 443端口timeout的错误

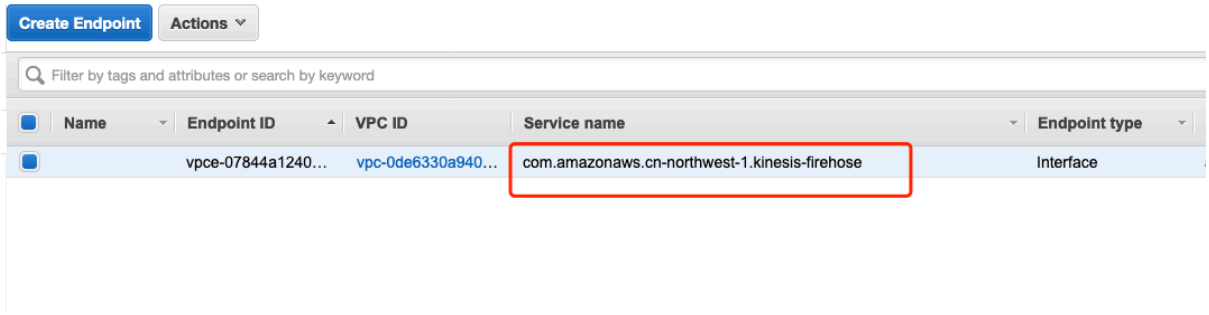

5.启动task

选择指定的vpc，subnet和sg。

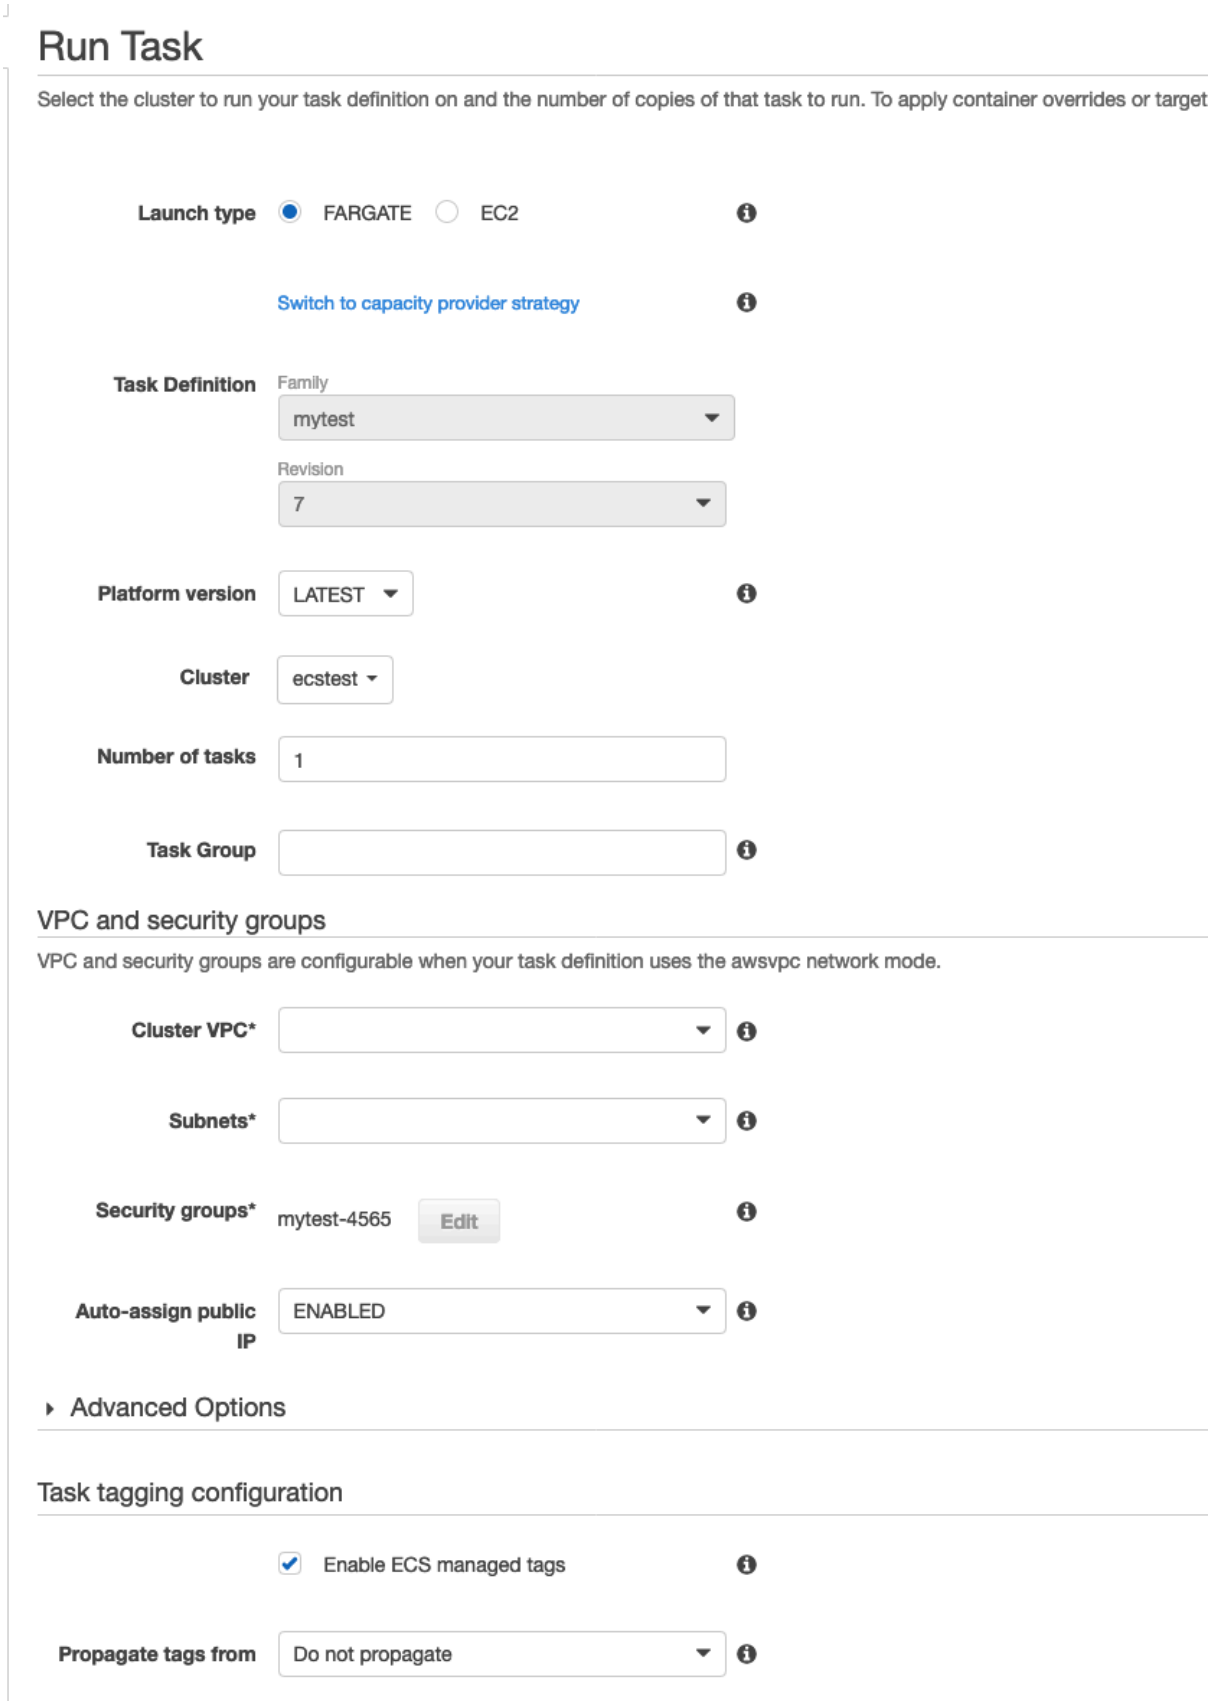

6. 检查

- fluentbit container在cloudwatch log中输出的日志：

No older events at this moment. *Retry*

2021-04-04T18:02:22.799+08:00	AWS for Fluent Bit Container Image Version 2.12.0
2021-04-04T18:02:22.840+08:00	[1mFluent Bit v1.7.2[0m
2021-04-04T18:02:22.840+08:00	* [1m[93mCopyright (C) 2019-2021 The Fluent Bit Authors[0m
2021-04-04T18:02:22.840+08:00	* [1m[93mCopyright (C) 2015-2018 Treasure Data[0m
2021-04-04T18:02:22.840+08:00	* Fluent Bit is a CNCF sub-project under the umbrella of Fluentd
2021-04-04T18:02:22.840+08:00	* https://fluentbit.io
2021-04-04T18:02:22.841+08:00	[2021/04/04 10:02:22] [info] [engine] started (pid=1)
2021-04-04T18:02:22.841+08:00	[2021/04/04 10:02:22] [info] [storage] version=1.1.1, initializing...
2021-04-04T18:02:22.841+08:00	[2021/04/04 10:02:22] [info] [storage] in-memory
2021-04-04T18:02:22.841+08:00	[2021/04/04 10:02:22] [info] [storage] normal synchronization mode, checksum disabled, max_chunks_up=128
2021-04-04T18:02:22.841+08:00	[2021/04/04 10:02:22] [info] [input:forward:forward.0] listening on unix:///var/run/fluent.sock
2021-04-04T18:02:22.841+08:00	[2021/04/04 10:02:22] [info] [input:forward:forward.1] listening on 127.0.0.1:24224
2021-04-04T18:02:22.841+08:00	[2021/04/04 10:02:22] [info] [input:tcp:tcp.2] listening on 127.0.0.1:8877
2021-04-04T18:02:22.842+08:00	time="2021-04-04T10:02:22Z" level=info msg="A new higher performance Firehose plugin has been released; you are using the old plugin. Check out the ne
2021-04-04T18:02:22.842+08:00	time="2021-04-04T10:02:22Z" level=info msg="[firehose 0] plugin parameter delivery_stream = 'fluentbit-access'"
2021-04-04T18:02:22.842+08:00	time="2021-04-04T10:02:22Z" level=info msg="[firehose 0] plugin parameter region = 'cn-northwest-1'"
2021-04-04T18:02:22.842+08:00	time="2021-04-04T10:02:22Z" level=info msg="[firehose 0] plugin parameter data_keys = ''"
2021-04-04T18:02:22.842+08:00	time="2021-04-04T10:02:22Z" level=info msg="[firehose 0] plugin parameter role_arn = ''"
2021-04-04T18:02:22.842+08:00	time="2021-04-04T10:02:22Z" level=info msg="[firehose 0] plugin parameter endpoint = ''"
2021-04-04T18:02:22.842+08:00	time="2021-04-04T10:02:22Z" level=info msg="[firehose 0] plugin parameter sts_endpoint = ''"
2021-04-04T18:02:22.842+08:00	time="2021-04-04T10:02:22Z" level=info msg="[firehose 0] plugin parameter time_key = ''"
2021-04-04T18:02:22.842+08:00	time="2021-04-04T10:02:22Z" level=info msg="[firehose 0] plugin parameter time_key_format = ''"
2021-04-04T18:02:22.842+08:00	time="2021-04-04T10:02:22Z" level=info msg="[firehose 0] plugin parameter log_key = ''"
2021-04-04T18:02:22.842+08:00	time="2021-04-04T10:02:22Z" level=info msg="[firehose 0] plugin parameter replace_dots = ''"
2021-04-04T18:02:22.842+08:00	time="2021-04-04T10:02:22Z" level=info msg="[firehose 0] plugin parameter simple_aggregation = 'false'"
2021-04-04T18:02:22.842+08:00	[2021/04/04 10:02:22] [info] [sp] stream processor started
2021-04-04T18:02:22.842+08:00	[2021/04/04 10:02:22] [info] [input:tail:tail.3] inotify_fs_add(): inode=262147 watch_fd=1 name=/data/access.log

- app container输出的日志