链滴

fluent-bit debug 调试，采集 kubernetes podIP

作者：liabio

原文链接：https://ld246.com/article/1635127076855

来源网站：链滴

有时候调试fluent-bit的配置，达到想要的输出效果，并不是件简单的事情，以下通过debug镜像调试fuent-bit采集kubernetes Pod的IP。

fluent-bit官方文档给出了用于调试的镜像：

https://docs.fluentbit.io/manual/installation/docker

dockerhub仓库链接为：

https://hub.docker.com/r/fluent/fluent-bit/

部署fluent-bit-debug

```
apiVersion: apps/v1
kind: Deployment
metadata:
  labels:
    app.kubernetes.io/name: fluent-bit-debug
  name: fluent-bit-debug
  namespace: kubesphere-logging-system
spec:
  progressDeadlineSeconds: 600
  replicas: 1
  revisionHistoryLimit: 10
  selector:
    matchLabels:
      app.kubernetes.io/name: fluent-bit-debug
  strategy:
    rollingUpdate:
      maxSurge: 25%
      maxUnavailable: 25%
    type: RollingUpdate
  template:
    metadata:
      creationTimestamp: null
      labels:
        app.kubernetes.io/name: fluent-bit-debug
      name: fluent-bit-debug
    spec:
      containers:
      - env:
        - name: NODE_NAME
          valueFrom:
            fieldRef:
              apiVersion: v1
              fieldPath: spec.nodeName
        command:
        - /usr/local/bin/sh
        - -c
        - sleep 9999
        image: fluent/fluent-bit:1.6.9-debug
        imagePullPolicy: IfNotPresent
        name: fluent-bit
        ports:
```

```
    - containerPort: 2020
      name: metrics
      protocol: TCP
    resources: {}
    terminationMessagePath: /dev/termination-log
    terminationMessagePolicy: File
    volumeMounts:
    - mountPath: /var/lib/docker/containers
      name: varlibcontainers
      readOnly: true
    - mountPath: /fluent-bit/config
      name: config
      readOnly: true
    - mountPath: /var/log/
      name: varlogs
      readOnly: true
    - mountPath: /var/log/journal
      name: systemd
      readOnly: true
    - mountPath: /fluent-bit/tail
      name: positions
  dnsPolicy: ClusterFirst
  restartPolicy: Always
  schedulerName: default-scheduler
  securityContext: {}
  serviceAccount: fluent-bit
  serviceAccountName: fluent-bit
  terminationGracePeriodSeconds: 30
  volumes:
  - hostPath:
      path: /var/lib/docker/containers
      type: ""
    name: varlibcontainers
  - name: config
    secret:
      defaultMode: 420
      secretName: fluent-bit-debug-config
  - hostPath:
      path: /var/log
      type: ""
    name: varlogs
  - hostPath:
      path: /var/log/journal
      type: ""
    name: systemd
  - emptyDir: {}
    name: positions
```

用到的secret：fluent-bit-debug-config如下，包含两个key：

原文链接：fluent-bit debug 调试，采集 kubernetes podIP

第一个parsers.conf：为空即可；配置细节可以参考官方文档：

Configuration File

第二个fluent-bit.conf配置需根据情况配置，以下主要在不同场景，给出fluent-bit.conf。也会涉及到ubesphere、Filter CRD一起使用。

fluent-bit使用

exec进入容器可以使用/fluent-bit/bin/fluent-bit调试：

```
/fluent-bit/bin #  ./fluent-bit -h
Usage: fluent-bit [OPTION]

Available Options
  -b  --storage_path=PATH       specify a storage buffering path
  -c  --config=FILE     specify an optional configuration file
  -d, --daemon          run Fluent Bit in background mode
  -f, --flush=SECONDS   flush timeout in seconds (default: 5)
  -F  --filter=FILTER     set a filter
  -i, --input=INPUT     set an input
  -m, --match=MATCH     set plugin match, same as '-p match=abc'
  -o, --output=OUTPUT   set an output
  -p, --prop="A=B"      set plugin configuration property
  -R, --parser=FILE     specify a parser configuration file
  -e, --plugin=FILE     load an external plugin (shared lib)
  -l, --log_file=FILE   write log info to a file
  -t, --tag=TAG         set plugin tag, same as '-p tag=abc'
  -T, --sp-task=SQL     define a stream processor task
  -v, --verbose         increase logging verbosity (default: info)
  -H, --http            enable monitoring HTTP server
  -P, --port            set HTTP server TCP port (default: 2020)
  -s, --coro_stack_size Set coroutines stack size in bytes (default: 24576)
  -q, --quiet           quiet mode
  -S, --sosreport       support report for Enterprise customers
  -V, --version         show version number
  -h, --help            print this help

Inputs
  cpu                 CPU Usage
  mem                 Memory Usage
  thermal             Thermal
  kmsg                Kernel Log Buffer
  proc                Check Process health
  disk                Diskstats
  systemd             Systemd (Journal) reader
  netif               Network Interface Usage
  docker              Docker containers metrics
  docker_events       Docker events
  tail                Tail files
  dummy               Generate dummy data
  head                Head Input
  health              Check TCP server health
  collectd            collectd input plugin
```

```
  statsd             StatsD input plugin
  serial             Serial input
  stdin              Standard Input
  syslog             Syslog
  exec               Exec Input
  tcp                TCP
  mqtt               MQTT, listen for Publish messages
  forward            Fluentd in-forward
  random             Random

Filters
  alter_size         Alter incoming chunk size
  aws                Add AWS Metadata
  record_modifier    modify record
  throttle           Throttle messages using sliding window algorithm
  kubernetes         Filter to append Kubernetes metadata
  modify             modify records by applying rules
  nest               nest events by specified field values
  parser             Parse events
  expect             Validate expected keys and values
  grep               grep events by specified field values
  rewrite_tag        Rewrite records tags
  lua                Lua Scripting Filter
  stdout             Filter events to STDOUT

Outputs
  azure              Send events to Azure HTTP Event Collector
  azure_blob         Azure Blob Storage
  bigquery           Send events to BigQuery via streaming insert
  counter            Records counter
  datadog            Send events to DataDog HTTP Event Collector
  es                 Elasticsearch
  exit               Exit after a number of flushes (test purposes)
  file               Generate log file
  forward            Forward (Fluentd protocol)
  http               HTTP Output
  influxdb           InfluxDB Time Series
  logdna             LogDNA
  loki               Loki
  kafka              Kafka
  kafka-rest         Kafka REST Proxy
  nats               NATS Server
  nrlogs             New Relic
  null               Throws away events
  plot               Generate data file for GNU Plot
  pgsql              PostgreSQL
  slack              Send events to a Slack channel
  splunk             Send events to Splunk HTTP Event Collector
  stackdriver        Send events to Google Stackdriver Logging
  stdout             Prints events to STDOUT
  syslog             Syslog
  tcp                TCP Output
  td                 Treasure Data
  flowcounter        FlowCounter
```

```
  gelf                GELF Output
  cloudwatch_logs     Send logs to Amazon CloudWatch
  kinesis_firehose    Send logs to Amazon Kinesis Firehose
  s3                  Send to S3

Internal
 Event Loop  = epoll
 Build Flags =  FLB_HAVE_HTTP_CLIENT_DEBUG FLB_HAVE_PARSER FLB_HAVE_RECORD_ACCE
SOR FLB_HAVE_STREAM_PROCESSOR JSMN_PARENT_LINKS JSMN_STRICT FLB_HAVE_TLS FLB
HAVE_AWS FLB_HAVE_SIGNV4 FLB_HAVE_SQLDB FLB_HAVE_METRICS FLB_HAVE_HTTP_SERV
R FLB_HAVE_SYSTEMD FLB_HAVE_FORK FLB_HAVE_TIMESPEC_GET FLB_HAVE_GMTOFF FLB_H
VE_UNIX_SOCKET FLB_HAVE_PROXY_GO FLB_HAVE_JEMALLOC JEMALLOC_MANGLE FLB_HAV
_LIBBACKTRACE FLB_HAVE_REGEX FLB_HAVE_UTF8_ENCODER FLB_HAVE_LUAJIT FLB_HAVE_C
TLS FLB_HAVE_ACCEPT4 FLB_HAVE_INOTIFY
```

简单配置文件

以下用一个简单的配置文件采集calico-node-* pod的日志：

```
[Service]
    Parsers_File    parsers.conf
[Input]
    Name    tail
    Path    /var/log/containers/*_kube-system_calico-node-*.log
    Refresh_Interval    10
    Skip_Long_Lines    true
    DB    /fluent-bit/bin/pos.db
    DB.Sync    Normal
    Mem_Buf_Limit    5MB
    Parser    docker
    Tag    kube.*
[Filter]
    Name    kubernetes
    Match    kube.*
    Kube_URL    https://kubernetes.default.svc:443
    Kube_CA_File    /var/run/secrets/kubernetes.io/serviceaccount/ca.crt
    Kube_Token_File    /var/run/secrets/kubernetes.io/serviceaccount/token
    Labels    false
    Annotations    true
[Output]
    Name    stdout
    Match_Regex    (?:kube|service)\.(.*)
```

fluent-bit-debug容器内使用以下命令启动测试：

```
/fluent-bit/bin # ./fluent-bit -c /fluent-bit/config/fluent-bit.conf
```

可以看到stdout日志输出：

```
[0] kube.var.log.containers.calico-node-lp4lm_kube-system_calico-node-cca502a39695f7452f
999af97bfbca5d74d2a372d94e0cacf2045f5f9721a81.log: [1634870260.700108403, {"log"=>"{"
og":"2021-10-22 02:37:40.699 [INFO][85] monitor-addresses/startup.go 774: Using autodetec
ed IPv4 address on interface bond4: 172.24.248.50/30\n","stream":"stdout","time":"2021-10-2
T02:37:40.700056471Z"}", "kubernetes"=>{"pod_name"=>"calico-node-lp4lm", "namespace_
```

ame"=>"kube-system", "pod_id"=>"5a829979-9830-4b9c-a3cb-eeb6eee38bdd", "annotation
"=>{"kubectl.kubernetes.io/restartedAt"=>"2021-10-20T23:00:27+08:00"}, "host"=>"node02",
"container_name"=>"calico-node", "docker_id"=>"cca502a39695f7452fd999af97bfbca5d74d
a372d94e0cacf2045f5f9721a81", "container_hash"=>"calico/node@sha256:bc4a631d553b38f
c169ea4cb8027fa894a656e80d68d513359a4b9d46836b55", "container_image"=>"calico/nod
:v3.19.1"}}]

截取重要部分，可以看到没经过处理采集到的k8s日志格式。


```
[
  {
    "kubernetes"=>{
      "pod_name"=>"calico-node-lp4lm",
      "namespace_name"=>"kube-system",
      "pod_id"=>"5a829979-9830-4b9c-a3cb-eeb6eee38bdd",
      "annotations"=>{
        "kubectl.kubernetes.io/restartedAt"=>"2021-10-20T23:00:27+08:00"
      },
      "host"=>"node02",
      "container_name"=>"calico-node",
      "docker_id"=>"cca502a39695f7452fd999af97bfbca5d74d2a372d94e0cacf2045f5f9721
81",
      "container_hash"=>"calico/node@sha256:bc4a631d553b38fdc169ea4cb8027fa894a6
6e80d68d513359a4b9d46836b55",
      "container_image"=>"calico/node:v3.19.1"
    }
  }
]
```


增加nest Filter

将kubernetes块展开，并添加kubernetes_前缀：

```
[Filter]
  Name    nest
  Match    kube.*
  Operation    lift
  Nested_under    kubernetes
  Add_prefix    kubernetes_
```

这次测试输出，截取重要部分：


```
{
  "kubernetes_pod_name"=>"calico-node-lp4lm",
  "kubernetes_namespace_name"=>"kube-system",
  "kubernetes_pod_id"=>"5a829979-9830-4b9c-a3cb-eeb6eee38bdd",
  "kubernetes_annotations"=>{
    "kubectl.kubernetes.io/restartedAt"=>"2021-10-20T23:00:27+08:00"
  },
  "kubernetes_host"=>"node02",
  "kubernetes_container_name"=>"calico-node",
  "kubernetes_docker_id"=>"cca502a39695f7452fd999af97bfbca5d74d2a372d94e0cacf2045
5f9721a81",
  "kubernetes_container_hash"=>"calico/node@sha256:bc4a631d553b38fdc169ea4cb8027f
```

```
894a656e80d68d513359a4b9d46836b55",
    "kubernetes_container_image"=>"calico/node:v3.19.1"
}
```

移除掉kubernetes_annotations块

```
[Filter]
    Name    modify
    Match    kube.*
    Remove    kubernetes_annotations
```

移除掉kubernetes_annotations块中的某字段

```
[Filter]
    Name    nest
    Match    kube.*
    Operation    lift
    Nested_under    kubernetes_annotations
    Add_prefix    kubernetes_annotations_
[Filter]
    Name    modify
    Match    kube.*
    Remove    kubernetes_annotations_kubectl.kubernetes.io/restartedAt
```

或者用正则：

```
[Filter]
    Name    nest
    Match    kube.*
    Operation    lift
    Nested_under    kubernetes_annotations
    Add_prefix    kubernetes_annotations_
[Filter]
    Name    modify
    Match    kube.*
    Remove_regex    kubernetes_annotations_kubectl*
```

修改kubernetes_annotations块中key名称

```
[Filter]
    Name    nest
    Match    kube.*
    Operation    lift
    Nested_under    kubernetes_annotations
    Add_prefix    kubernetes_annotations_
[Filter]
    Name    modify
    Match    kube.*
    Rename    kubernetes_annotations_kubectl.kubernetes.io/restartedAt podIPs
```

修改之后：

```
[
```

```
    {
        "kubernetes_pod_name"=>"calico-node-lp4lm",
        "kubernetes_namespace_name"=>"kube-system",
        "kubernetes_pod_id"=>"5a829979-9830-4b9c-a3cb-eeb6eee38bdd",
        "kubernetes_host"=>"node02",
        "kubernetes_container_name"=>"calico-node",
        "kubernetes_docker_id"=>"cca502a39695f7452fd999af97bfbca5d74d2a372d94e0cacf205f5f9721a81",
        "kubernetes_container_hash"=>"calico/node@sha256:bc4a631d553b38fdc169ea4cb802fa894a656e80d68d513359a4b9d46836b55",
        "kubernetes_container_image"=>"calico/node:v3.19.1",
        "podIPs"=>"2021-10-20T23:00:27+08:00"
    }
]
```

结合ks配置采集podIPs

结合kubesphere Filter CR，配置采集podIPs，并去掉其他不相关的annotations。

因使用calico作为CNI，所以在pod annotations中会被添加上podIP相关的注解。

需要保留注解中的某一个key（cni.projectcalico.org/podIPs），移除掉其他key，所以下面将要保的key修改名称之后，移除掉整个annotations。

kubernetes Filter CR配置如下：

```
apiVersion: logging.kubesphere.io/v1alpha2
kind: Filter
metadata:
  labels:
    logging.kubesphere.io/component: logging
    logging.kubesphere.io/enabled: 'true'
  name: kubernetes
  namespace: kubesphere-logging-system
spec:
  filters:
    - kubernetes:
        annotations: true
        kubeCAFile: /var/run/secrets/kubernetes.io/serviceaccount/ca.crt
        kubeTokenFile: /var/run/secrets/kubernetes.io/serviceaccount/token
        kubeURL: 'https://kubernetes.default.svc:443'
        labels: false
    - nest:
        addPrefix: kubernetes_
        nestedUnder: kubernetes
        operation: lift
    - nest:
        addPrefix: kubernetes_annotations_
        nestedUnder: kubernetes_annotations
        operation: lift
    - modify:
        rules:
          - remove: stream
          - remove: kubernetes_pod_id
```

```
      - remove: kubernetes_host
      - remove: kubernetes_container_hash
      - rename:
          kubernetes_annotations_cni.projectcalico.org/podIPs: kubernetes_podIPs
      - removeRegex: kubernetes_annotations*
  - nest:
      nestUnder: kubernetes_annotations
      operation: nest
      removePrefix: kubernetes_annotations_
      wildcard:
        - kubernetes_annotations_*
  - nest:
      nestUnder: kubernetes
      operation: nest
      removePrefix: kubernetes_
      wildcard:
        - kubernetes_*
match: kube.*
```

由kubernetes Filter CR生成的fluent-bit config配置如下（只看Filter部分，Input、Output CR被略）

```
[Service]
    Parsers_File    parsers.conf
[Input]
    Name    tail
    Path    /var/log/containers/*.log
    Exclude_Path    /var/log/containers/*_kubesphere-logging-system_events-exporter*.log,/va
/log/containers/kube-auditing-webhook*_kubesphere-logging-system_kube-auditing-webho
k*.log
    Refresh_Interval    10
    Skip_Long_Lines    true
    DB    /fluent-bit/tail/pos.db
    DB.Sync    Normal
    Mem_Buf_Limit    5MB
    Parser    docker
    Tag    kube.*
[Filter]
    Name    kubernetes
    Match    kube.*
    Kube_URL    https://kubernetes.default.svc:443
    Kube_CA_File    /var/run/secrets/kubernetes.io/serviceaccount/ca.crt
    Kube_Token_File    /var/run/secrets/kubernetes.io/serviceaccount/token
    Labels    false
    Annotations    true
[Filter]
    Name    nest
    Match    kube.*
    Operation    lift
    Nested_under    kubernetes
    Add_prefix    kubernetes_
[Filter]
    Name    nest
```

```
    Match    kube.*
    Operation    lift
    Nested_under    kubernetes_annotations
    Add_prefix    kubernetes_annotations_
[Filter]
    Name    modify
    Match    kube.*
    Remove    stream
    Remove    kubernetes_pod_id
    Remove    kubernetes_host
    Remove    kubernetes_container_hash
    Rename    kubernetes_annotations_cni.projectcalico.org/podIPs    kubernetes_podIPs
    Remove_regex    kubernetes_annotations*
[Filter]
    Name    nest
    Match    kube.*
    Operation    nest
    Wildcard    kubernetes_annotations_*
    Nest_under    kubernetes_annotations
    Remove_prefix    kubernetes_annotations_
[Filter]
    Name    nest
    Match    kube.*
    Operation    nest
    Wildcard    kubernetes_*
    Nest_under    kubernetes
    Remove_prefix    kubernetes_
[Output]
    Name    es
    Match_Regex    (?:kube|service)\.(.*)
    Host    es-cdc-a-es-http.cdc.svc.xke.test.cn
    Port    9200
    HTTP_User    elastic
    HTTP_Passwd    elasticpwd
    Logstash_Format    true
    Logstash_Prefix    ks-logstash-log
    Time_Key    @timestamp
```

可以在kibana看到采集到的podIP：

原文链接：fluent-bit debug 调试，采集 kubernetes podIP