



# Method to Ship Docker, Journalctl, and QNS Logs to External EFK Stack

---

- [Feature Description, on page 1](#)
- [Configuration to Fetch Journalctl, on page 1](#)
- [Configuration to fetch the consolidated-qns logs and mongo logs, on page 2](#)
- [Configuration for local Log forwarding, on page 2](#)
- [Configuration for Controlling the Interval and Size Forwarding , on page 3](#)
- [Configuration to Forward Remote Logs , on page 3](#)
- [Configuration for Log Filtration, on page 4](#)

## Feature Description

vDRA supports a unified method to forward all required logs such as journalctl, consolidated-qns logs, mongo logs to elasticsearch. You can have a consolidated view of all the logs with Elasticsearch Fluentbit Kibana (EFK) stack. In addition, using Kibana you can visualize and filter required logs for analysis.

Elasticsearch is an open source, full-text search and analytics engine, based on the Apache Lucene search engine. Elasticsearch indexes and stores the data.

Fluent Bit is an open source Log Processor and Forwarder which allows you to collect any data like metrics and logs from different sources, enrich them with filters and send them to multiple destinations. Fluent-Bit takes care of data collection and processing.

Kibana is a visualization layer that works on top of Elasticsearch, providing users with the ability to analyze and visualize the data. Kibana provides a user interface for querying the data and visualizing.

Logs are collected within each VM and same are forwarded to one of the OAM VMs. The logs are then forwarded from the corresponding OAM VM to external servers. Logs can be filtered based on keywords before it is sent to the elastic search server. For more information about CLI Command configurations, see the *CLI Commands* chapter in the *CPS vDRA Operations Guide*.

## Configuration to Fetch Journalctl

Logs from journald are available to fluent-bit through the input plugin Systemd. With this plugin, Journalctl logs are available with required journald key value pairs. The configuration file is available at every VM at `/etc/td-agent-bit/td-agent-bit.conf`. For example:

```
[INPUT]
  Name systemd
  Tag host.*
  Systemd_Filter _SYSTEMD_UNIT=docker.service <optional>
```

### Limitations

Currently the journalctl logs are available as single line entries in fluent-bit. Multiline parsing is not available for trace errors in logs.

## Configuration to fetch the consolidated-qns logs and mongo logs

consolidated-qns logs and mongo logs are part of the docker logs mounted to the host. These logs are available as part of the tail plugin, which is then forwarded to the required OAM vms.

### consolidated-qns logs:

Configuration files for fetching the consolidated-qns logs are available at `/etc/td-agent-bit/td-agent-bit.conf` for each VMs control in DRA vnf. For example:

```
[INPUT]
  name tail
  path /data/cc-monitor-s102/var/log/broadhop/consolidated-qns.log
  tag consolidated_qns_logs
```

### mongo logs:

Configuration files for fetching the mongo logs are available at `/etc/td-agent-bit/td-agent-bit.conf` for each VM in Database Base (DB) vnf. For example:

```
[INPUT]
  name tail
  path /data/mongod-node/db/mongo-*.log
  tag logs_on_db
```

## Configuration for local Log forwarding

You can enable Log forwarding locally from all servers in DRA to forward the local logs to one of the control vm from where the logs can be extracted to external servers.

Enable Configuration for local forwarding:

Configuration file: `/etc/td-agent-bit/td-agent-bit.conf`

Sample Configuration:

```
[OUTPUT]
  name forward
  match *
  host <OAM-VIP>
  port 24224
```

Make sure to configure the OAM VIP using CLI commands.

## Configuration for Controlling the Interval and Size Forwarding

Use the following configuration for controlling the interval when logs are forwarded and the size that can be forwarded with each batch. The following configuration is available on the OAM vm at `/etc/td-agent-bit/td-agent-bit.conf`

Sample Configuration:

```
[SERVICE]
# Flush
# set an interval of seconds before to flush records to a destination
flush          400
storage.path   /var/log/flb/
storage.sync   normal
storage.checksum off
storage.max_chunks_up 80
storage.backlog.mem_limit 500M

[INPUT]
name          forward
listen       0.0.0.0
port         24224
storage.type  filesystem
```

The “flush” interval decides the interval at which logs are flushed to output.

## Configuration to Forward Remote Logs

Enable the Remote log forwarding through elasticsearch plugin of Fluent-bit output configuration. Configurations for elasticsearch configuration with fluent-bit are available at `/etc/td-agent-bit/td-agent-bit.conf` on proxy vm (OAM).

This enables the proxy vm (control vm in this case) to forward all the logs collected to the external server.

Sample Configuration:

```
[OUTPUT]
name es
match *
host <172.18.63.228>
port <9200>
index fluent_bit
HTTP_User <username>
HTTP_Passwd <password>
Logstash_Format on
Retry_Limit 5
```

The host IP (elasticsearch IP), port, username, password are configurable with CLI commands

Password authentication is enabled for external server with the elasticsearch plugin of fluent-bit.

## Monitoring Healthcheck of Elasticsearch Server

The ElasticSearch Server (External Server) (elasticsearch server) is monitored to DRA if its reachable to the OAM vms. If there is an unreachability, the alert is triggered. If there is no configuration provided for external

server IP, no alert is observed. Use the **ELASTICSEARCH\_NOT\_REACHABLE** alert to notify user if the External elasticsearch server is reachable to DRA. If this is not reachable to DRA, an alert is raised.

## Configuration for Log Filtration

Use the following sample configuration for filtering the logs before it is forwarded to the external servers. The filter section configured at `/etc/td-agent-bit/td-agent-bit.conf` for the OAM vm:

```
[FILTER]
name grep
match *
regex log aa
```

Configure the "regex" through CLI to apply the pattern and filter logs before it is sent to the external server.