Lab 1: Filebeat

1: install filebeat , if it is already there remove

dpkg —purge filebeat

Make sure delete /etc/filebeat , /usr/share/filebeat and /var/lib/filebeat

2: Go to /etc/filebeat/filebeat.yml and open in editor

Do following changes.

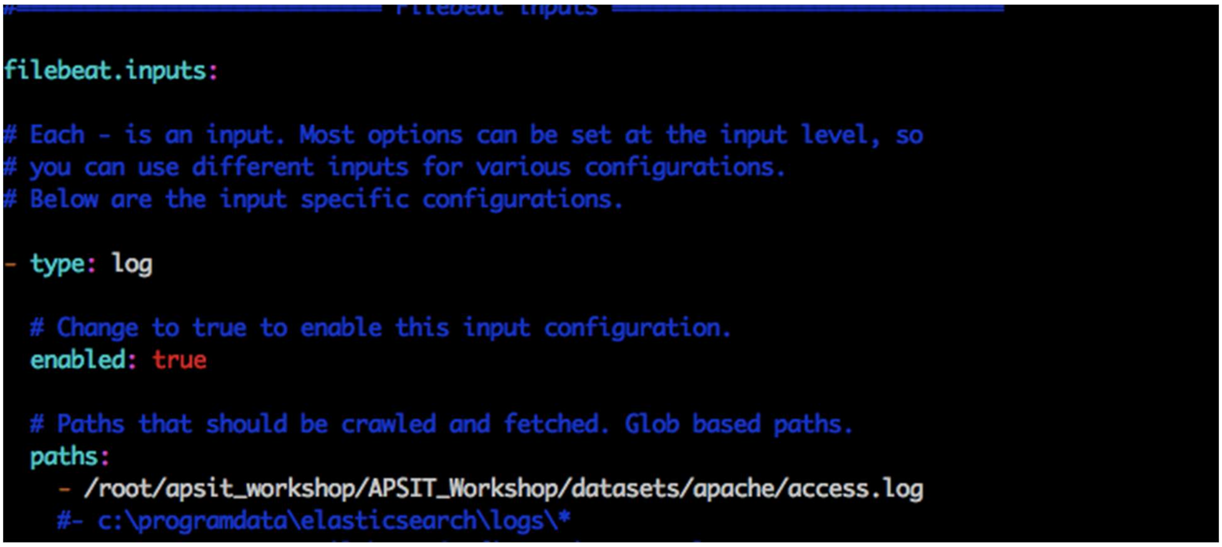

```
#============================= Filebeat inputs =============================

filebeat.inputs:

# Each - is an input. Most options can be set at the input level, so
# you can use different inputs for various configurations.
# Below are the input specific configurations.

- type: log

  # Change to true to enable this input configuration.
  enabled: true

  # Paths that should be crawled and fetched. Glob based paths.
  paths:
    - /root/apsit_workshop/APSIT_Workshop/datasets/apache/access.log
    #- c:\programdata\elasticsearch\logs\*
```

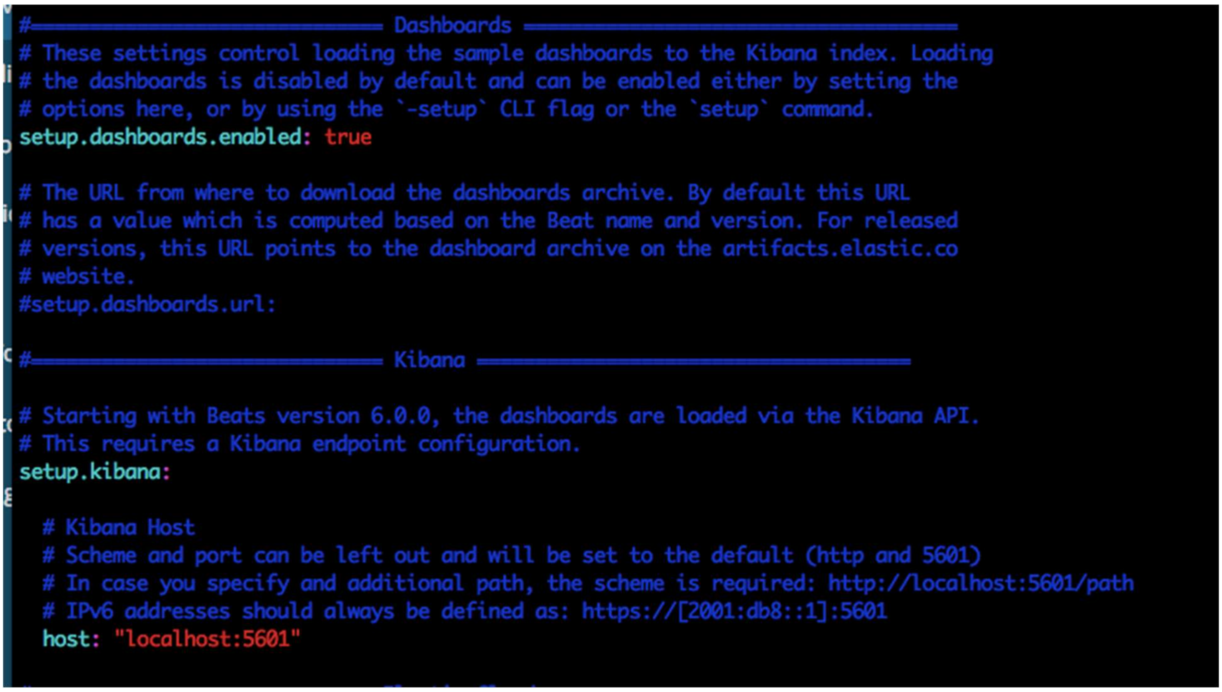

```
#============================== Dashboards =====================================
# These settings control loading the sample dashboards to the Kibana index. Loading
# the dashboards is disabled by default and can be enabled either by setting the
# options here, or by using the `-setup` CLI flag or the `setup` command.
setup.dashboards.enabled: true

# The URL from where to download the dashboards archive. By default this URL
# has a value which is computed based on the Beat name and version. For released
# versions, this URL points to the dashboard archive on the artifacts.elastic.co
# website.
#setup.dashboards.url:

#============================== Kibana =====================================

# Starting with Beats version 6.0.0, the dashboards are loaded via the Kibana API.
# This requires a Kibana endpoint configuration.
setup.kibana:

  # Kibana Host
  # Scheme and port can be left out and will be set to the default (http and 5601)
  # In case you specify and additional path, the scheme is required: http://localhost:5601/path
  # IPv6 addresses should always be defined as: https://[2001:db8::1]:5601
  host: "localhost:5601"
```

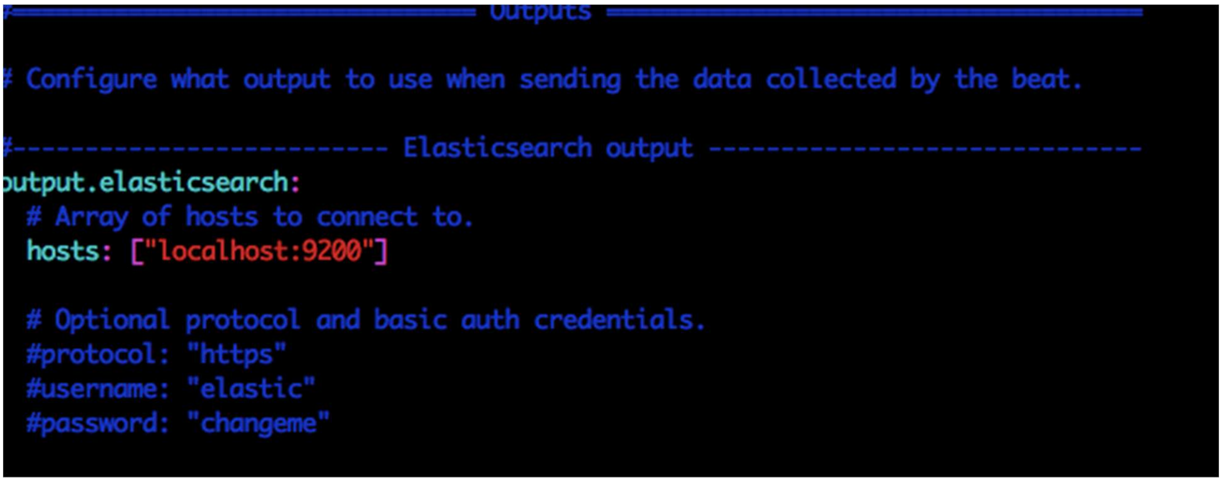

```
#================================ Outputs =====================================

# Configure what output to use when sending the data collected by the beat.

#-------------------------- Elasticsearch output ------------------------------
output.elasticsearch:
  # Array of hosts to connect to.
  hosts: ["localhost:9200"]

  # Optional protocol and basic auth credentials.
  #protocol: "https"
  #username: "elastic"
  #password: "changeme"
```

save the changes

3: start the filbeat service

service filebeat start

4: Go to kibana and check Index pattern and it should appear filebat-* .
Select the same and create the index

5: Go to discover and observe the logs
Total documents are 300,000

Lab 2: Logstash :

1: Install logstash

dpkg –i Logstash

2: go to /etc/Logstash/conf.d and create the access log filter module.

Copy and past the given access_log_1.conf file in conf.d folder

3: Test the given file for configuration OK.

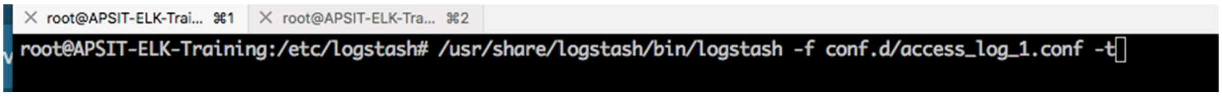

```
root@APSIT-ELK-Training:/etc/logstash# /usr/share/logstash/bin/logstash -f conf.d/access_log_1.conf -t
```

4: Now run the given file

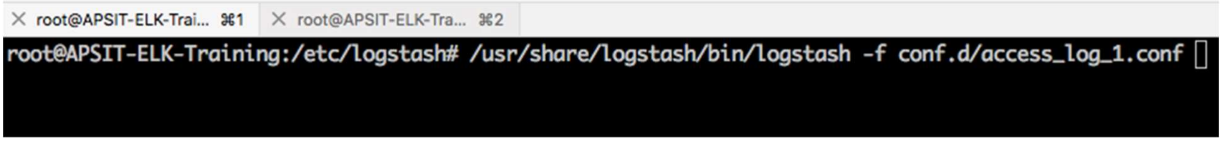

```
root@APSIT-ELK-Training:/etc/logstash# /usr/share/logstash/bin/logstash -f conf.d/access_log_1.conf
```

5: Go to Kibana and Create the Index pattern Logstash-*

6: Now go to discover page and select the Logstash index and check the data , it should show approx. ~300,000 documents (check the date 2 years).

7: Load the dashboard template into kibana

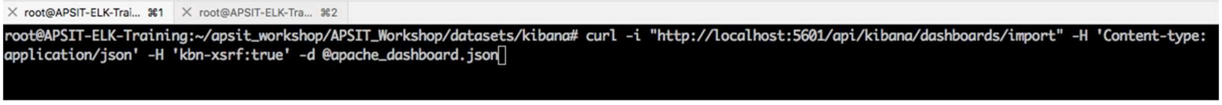

```
root@APSIT-ELK-Training:~/apsit_workshop/APSIT_Workshop/datasets/kibana# curl -i "http://localhost:5601/api/kibana/dashboards/import" -H 'Content-type:
application/json' -H 'kbn-xsrf:true' -d @apache_dashboard.json
```

8: Now go to Dashboard and slect Apache Access log and observe the dashboard.

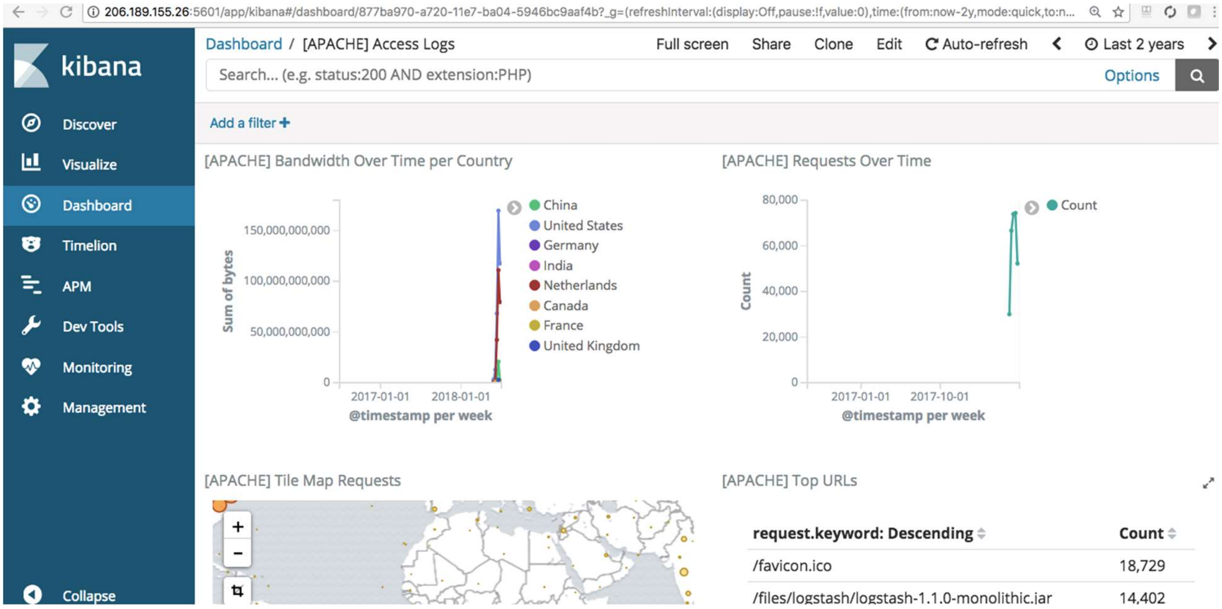

request.keyword: Descending	Count
/favicon.ico	18,729
/files/logstash/logstash-1.1.0-monolithic.jar	14,402