

ANNEXURE - 2  
Oracle Banking Branch  
Release 14.5.0.0.0  
Part Number F41897-01  
May 2021



---

# Table of Contents

- 1. ANNEXURE - 2 ..... 1-1**
  - 1.1 INTRODUCTION ..... 1-1
- 2. DOCUMENT TRACING ZIPKIN ..... 2-1**
  - 2.1 INSTALLATION OF ZIPKIN..... 2-1
    - 2.1.1 Download and Running ..... 2-1
  - 2.2 ZIPKIN USER INTERFACE..... 2-1
- 3. MONITORING ELK ..... 3-1**
  - 3.1 INTRODUCTION ..... 3-1
  - 3.2 ARCHITECTURE..... 3-1
  - 3.3 INSTALLING AND CONFIGURING ELK ..... 3-2
    - 3.3.1 Setup ..... 3-2

### 1.1 Introduction

This guide is a supporting document for the installation of Zipkin and ELK. You can find the reference in the respective installation guides.

---

## 2. Document Tracing Zipkin

### 2.1 Installation of Zipkin

You can download and run the application to install Zipkin.

#### 2.1.1 Download and Running

Zipkin works as an independent application and it can be downloaded as a runnable jar from the official website of Zipkin <https://zipkin.io/>. The latest version of Zipkin needs a Java version above 8.

The direct download link of jar is as follows:

[https://search.maven.org/remote\\_content?g=io.zipkin&a=zipkin-server&v=LATEST&c=exec](https://search.maven.org/remote_content?g=io.zipkin&a=zipkin-server&v=LATEST&c=exec)

The downloaded jar can be executed using the java -jar JAR\_NAME command.

The configuration of Zipkin can be done environment variables. The port of the Zipkin can be set using QUERY\_PORT environment variable.

The application starts on the port number assigned for QUERY\_PORT environment variable or its default value of 9411. The web UI of Zipkin can be accessed at <http://localhost:PORT>.

### 2.2 Zipkin User Interface

The basic layout of Zipkin looks as follows:

The screenshot displays the Zipkin web user interface. At the top, there are navigation links: "Investigate system behavior", "Find a trace", "View Saved Trace", and "Dependencies". On the right, there are buttons for "Try Lens UI", "Go to trace", and "Search". The main search area contains several filters: "Service Name" (set to "zipkin"), "Span Name" (set to "all"), "Remote Service Name" (set to "all"), and "Lookback" (set to "15 minutes"). Below these are "Annotation Query" (with an example: "http.path=/foo/bar/ and cluster=foo and cache.miss"), "Duration (µs) >=" (with an example: "100ms or 5s"), "Limit" (set to "10"), and "Sort" (set to "Longest First"). A blue "Find Traces" button is located below the filters. A light blue banner at the bottom of the search area contains the text: "Please select the criteria for your trace lookup."

We can find the traces of required API calls and services using the above search options given in the user interface. The search options given in the user interface are self-explanatory and there is another UI option (Try Lens UI). It is given a different user interface with same functionality.

The list of the traces can be seen like the above screen. Some error API calls are made to showcase how to track errors. The blue listings show the successful API hits and the red listings indicate errors. Each block indicates a single trace in the listings.

Opening an individual trace shows the below shown screen.

Services	2.163s	432.639ms	865.278ms	1.298s	1.731s	2.163s
zipkin	2.163s : http/api1	-	-	-	-	-
zipkin	-	1.001s : api1	-	0	-	-
zipkin	-	-	-	1.068s : http/api2	-	-
zipkin	-	-	-	1.001s : api2	-	-

The above shown image describes the time taken for each block. There are 2 custom spans created inside 2 service calls, so there are total of 4 blocks. The time taken for individual block can be seen above. Clicking an individual block shows the following details.

Date Time	Relative Time	Annotation	Address
9/10/2019, 4:11:23 PM		Server Start	
9/10/2019, 4:11:25 PM	2.163s	Server Finish	

Key	Value
http.host	localhost
http.method	GET
http.path	/api1
http.status_code	200
http.url	http://localhost:8080/api1
mvc.controller.class	Controller
mvc.controller.method	api1
spring.instance_id	eswarperabathini.in.oracle.com:Zipkin

traceld	9d63642d72ab6f9f
spanid	9d63642d72ab6f9f

The details of the specific span block are shown above and the logging events can also be seen in the Zipkin UI as small circular blocks. An example of error log is shown below:

Duration: 1.026s    Services: 1    Depth: 2    Total Spans: 3

Expand All    Collapse All

zipkin x3

Services	Start Time	End Time	Duration
zipkin	-1.026s	0s	1.026s
zipkin	-1.001s	-0.801s	200.134ms
zipkin	-0.801s	-0.581s	220.267ms
zipkin	-0.581s	-0.361s	220.401ms
zipkin	-0.361s	-0.141s	220.534ms
zipkin	-0.141s	0s	160.8ms

Clicking the **Error** portion gives the clear detail about the error and where the error has arised. AN example is shown below:

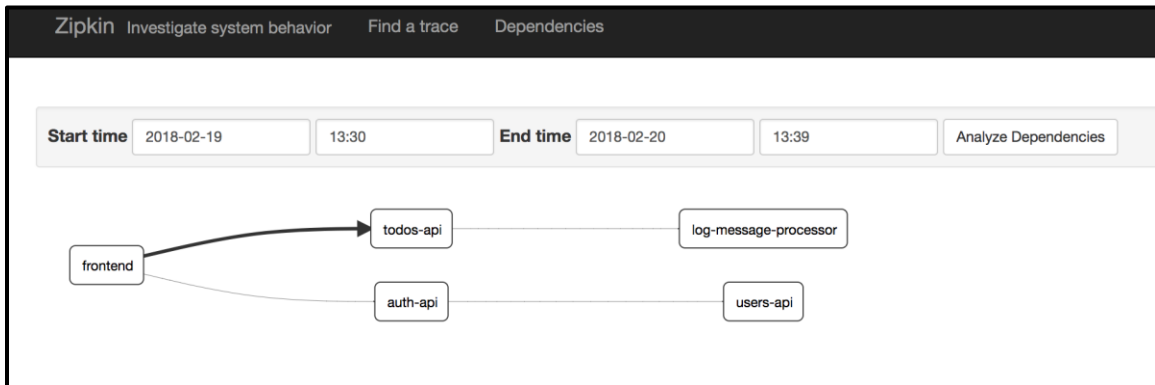
Services: zipkin

Date Time	Relative Time	Annotation	Address
9/11/2019, 6:09:01 PM		Server Start	
9/11/2019, 6:09:02 PM	1.026s	Server Finish	

Key	Value
error	Request processing failed; nested exception is org.springframework.web.client.HttpServerErrorException: 500 null
http.host	localhost
http.method	GET
http.path	/api1
http.status_code	500
http.url	http://localhost:8080/api1
mvc.controller.class	BasicErrorController
mvc.controller.method	errorHtml
spring.instance_id	eswarperabathini.in.oracle.com:Zipkin

If the Lens UI is used in Zipkin, the above screen shots are not applicable, but are relatable to the Lens UI as well.

Traces of the application can be found using Traceld, which can be found in the debug logs of the deployment when spring-cloud-sleuth is included in the dependencies (Included in spring-cloud-starter-zipkin dependency). Clicking the **Dependencies** tab gives the dependency graph info between micro-services. An example dependency graph is shown below:



## 3. Monitoring ELK

### 3.1 Introduction

ELK Stack was a collection of the following open-source products:


- Elasticsearch
- Logstash
- Kibana

Elasticsearch is an open source, full-text search and analysis engine, based on the Apache Lucene search engine. Logstash is a log aggregator that collects data from various input sources, executes different transformations and enhancements and then ships the data to various supported output destinations. Kibana is a visualization layer that works on top of Elasticsearch, providing users with the ability to analyze and visualize the data.

Together, these different components are most commonly used for monitoring, troubleshooting, and securing IT environments. Logstash take care of data collection and processing, Elasticsearch indexes and stores the data, and Kibana provides a user interface for querying the data and visualizing it.

### 3.2 Architecture

The below architecture provides a comprehensive solution for handling all the required facets:



Spring cloud Sleuth also provides additional functionality to keep trace of the application calls by providing us a way to create intermediate logging events. Thus, the Spring Cloud Sleuth dependency must be added to applications.



## 3.3 Installing and Configuring ELK

To install and configure ELK Stack, make sure the versions of the 3 software are same. Download the latest version of the following:

- Logstash
- Elastic Search
- Kibana

The installation guides are given below.

- Logstash : <https://www.elastic.co/guide/en/logstash/current/installing-logstash.html>
- Elastic Search : <https://www.elastic.co/guide/en/elasticsearch/reference/current/install-elasticsearch.html>
- Kibana : <https://www.elastic.co/guide/en/kibana/current/install.html>

Follow the process as given in the following sub-sections, after completing the download process of ELK.

### 3.3.1 Setup

The setup includes the following steps:

- Start Elastic Search
- Setup Logstash and Start
- Setup Kibana and Start

#### 3.3.1.1 Start ElasticSearch

1. Go to Elasticsearch root folder and use nohup to start the Elasticsearch process as below:

```
> nohup ./bin/elasticsearch
```

#### 3.3.1.2 Setup Logstash and Start

1. Create a new **logstash.conf** file that provides the required file parsing and integration to Elasticsearch.

**logstatsh.conf:**

```
#Point to the application logs
input {
  file {
    type => "java"
    path => "/scratch/app/work_area/app_Logs/*.Log"
    codec => multiline {
      pattern => "^[Y]{YEAR}-[M]{MONTHNUM}-[D]{MONTHDAY} [T]{TIME}.*"
      negate => "true"
      what => "previous"
    }
  }
}

#Provide the parsing logic to transform logs into JSON
filter {
```

```

#If log line contains tab character followed by 'at' then we
will tag that entry as stacktrace
if [message] =~ "\tat" {
  grok {
    match => ["message", "^(\\tat)"]
    add_tag => ["stacktrace"]
  }
}

#Grokking Spring Boot's default Log format
grok {
  match => [ "message",
            "(?<timestamp>{%YEAR}-{%MONTHNUM}-{%MONTHDAY}
            {%TIME}) {%LOGLEVEL:level} {%NUMBER:pid} --- \[(?<thread>[A-Za-
            z0-9-]+)\] [A-Za-z0-9.]*\.(?<class>[A-Za-z0-
            9#_]+)\s*:\s+(?<Logmessage>.*)",
            "message",
            "(?<timestamp>{%YEAR}-{%MONTHNUM}-{%MONTHDAY}
            {%TIME}) {%LOGLEVEL:level} {%NUMBER:pid} --- .+?
            :\s+(?<Logmessage>.*)"
          ]
}

# pattern matching logback pattern
grok {
  match =>
{ "message" => "%{TIMESTAMP_ISO8601:timestamp}\s+{%LOGLEVEL:seve
rity}\s+\[%{DATA:service},%{DATA:trace},%{DATA:span},%{DATA:expo
rtable}\]\s+\[%{DATA:environment}\]\s+\[%{DATA:tenant}\]\s+\[%{D
ATA:user}\]\s+\[%{DATA:branch}\]\s+{%DATA:pid}\s+---
\s+\[%{DATA:thread}\]\s+{%DATA:class}\s+:\s+{%GREEDYDATA:rest}"
}
}

#Parsing out timestamps which are in timestamp field thanks to
previous grok section
date {
  match => [ "timestamp" , "yyyy-MM-dd HH:mm:ss.SSS" ]
}
}

#Ingest Logs to Elasticsearch
output {
  elasticsearch { hosts => ["localhost:9200"] }
  stdout { codec => rubydebug }
}
}

```

2. Start Logstash process

```
>nohup ./bin/logstash -f logstash.conf
```

3.3.1.3 Setup Kibana and start

- 1. Navigate to the **kibana.yml** available under <kibana\_setup\_folder>/config and modify the file to include the below:

```
#Uncomment the below line and update the IP address to your host machine IP.
server.host: "xx.xxx.xxx.xx"
#Provide the elasticsearch url. If this is running on the same machine then you can use the below config as is
elasticsearch.url: "http://localhost:9200"
```

- 2. Start Kibana process using the below command:

```
>nohup ./bin/kibana
```

A view of the Kibana dashboard is given below:





## ANNEXURE - 2

May 2021

Version 14.5.0.0.0

Oracle Financial Services Software Limited  
Oracle Park  
Off Western Express Highway  
Goregaon (East)  
Mumbai, Maharashtra 400 063  
India

Worldwide Inquiries:

Phone: +91 22 6718 3000

Fax: +91 22 6718 3001

<https://www.oracle.com/industries/financial-services/index.html>

Copyright © 2021, Oracle and/or its affiliates. All rights reserved.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate failsafe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

This software or hardware and documentation may provide access to or information on content, products and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.