MONITORING MIXED-LANGUAGE APPLICATIONS WITH ELASTICSEARCH, LOGSTASH AND KIBANA (ELK)

A. De Dios Fuente, O. Ø. Andreassen, C. Charrondière, CERN, Geneva, Switzerland

2015 ICALEPCS melbourne • australia

ABSTRACT

Application logging and system diagnostics is nothing new. Ever since we had the first computers scientists and engineers have been storing information about their systems, making it easier to understand what is going on and, in case of failures, what went wrong. Unfortunately there are as many different standards as there are file formats, storage types, locations, operating systems, etc. Recent development in web technology and storage has made it much simpler to gather all the different information in one place and dynamically adapt the display. With the introduction of Logstash with Elasticsearch as a backend, we store, index and query data, making it possible to display and manipulate data in whatever form one wishes. With Kibana as a generic and modern web interface on top, the information can be adapted at will. In this paper we will show how we can process almost any type of structured or unstructured data source. We will also show how data can be visualised and customised on a per user basis and how the system scales when the data volume grows.

OUR CHALLENGE

In our current system, there is a **large variety of sources**; LabVIEW applications running on CompactRIO and PXI targets, Apache Tomcat servers, Java services, C++ applications and extensions running on the most popular operating systems (Linux, Windows and OS X). In such systems, most of the logs have a different or weak structure. The main requirements for the desired system are:

- Support multiple log formats.
- Support different communication and network protocols.
- Centralized data messages.
- Have a web-viewer to analyse the logs.
- Be able to get statistics of the stored data.

THE ELK STACK

Elasticsearch, Logstash and Kibana are three open sources projects which combined are known as the ELK stack.

Logstash is aimed to unify and normalise data from different sources in real time. It contains a rich collection of input and output plugins, for example:

- Logs: log4j for Java, syslog.
- Databases: Redis, SQLite, MongoDB.
- Network: UPD, TCP, WebSocket.
- Data streams: RabbitMQ, ZeroMQ.

Elasticsearch, as data store, is meant for handling real time data that needs to be processed and analysed in a rapid manner.

Kibana is the web interface to visualise and interact with the data through powerful graphics. Helps to understand large volumes of data and rapidly detect patterns or irregularities in them.

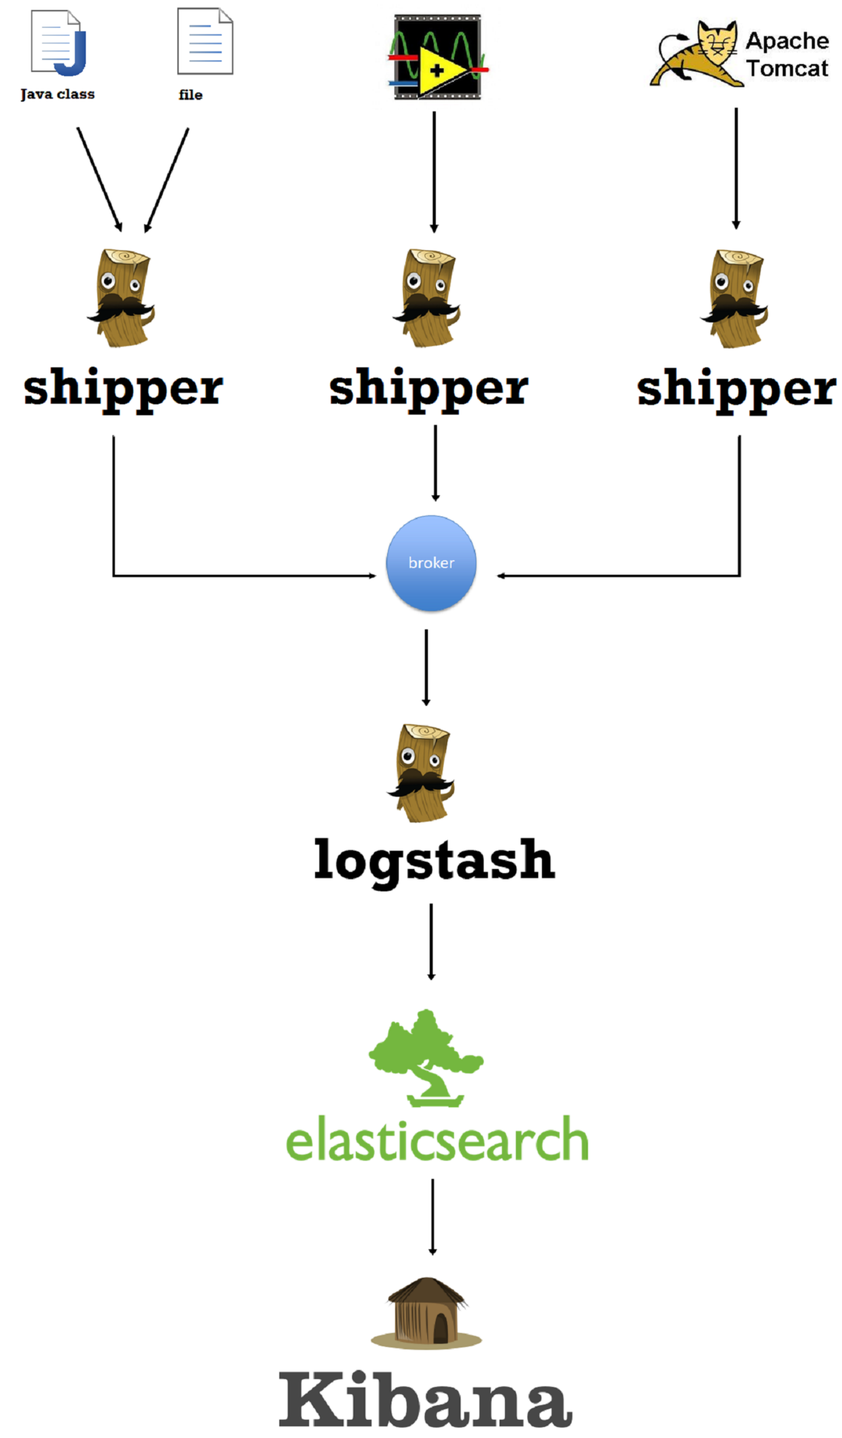

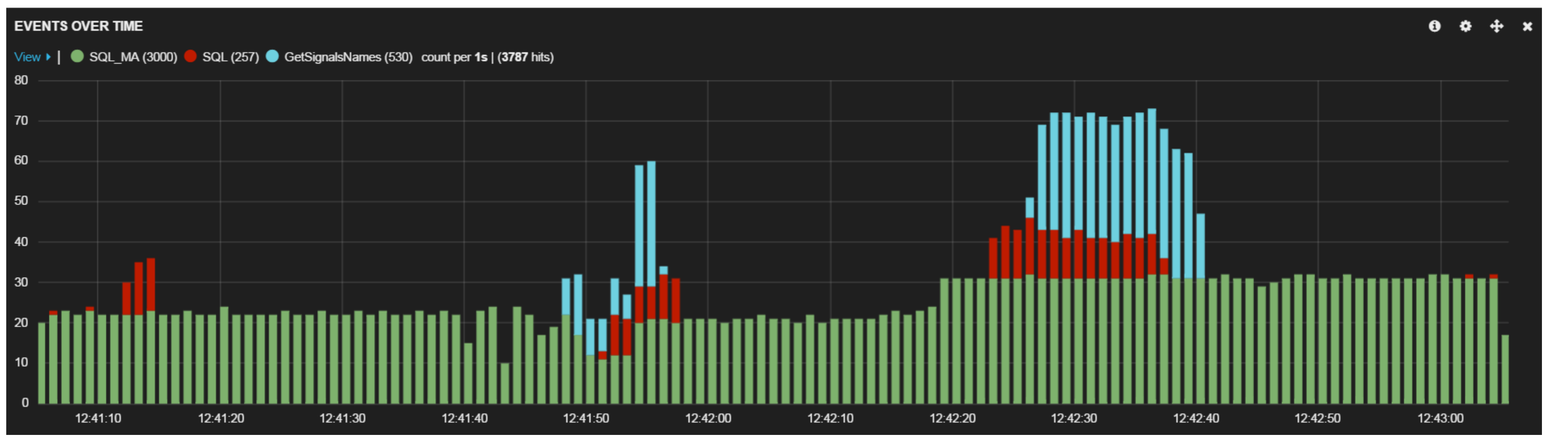

KIBANA DASHBOARD

The Kibana interface is simple and intuitive. There is no need to have a web development background to use it. Since the data and the purposes are different per user, each of them can customize a dashboard according to his needs.
Log messages are indexed and tagged by Logstash, then filtered with Kibana to gather metrics and statistics.

TESTS AND INTEGRATION

- Installation of the system was really straightforward; less than a day with the basic configuration.
- Throughput is 5000 messages per second.
- Developed a Logger tool in LabVIEW where messages are sent by UDP and the implementation took less than a day. It is included in the RADE framework where all LabVIEW users can benefit from it.

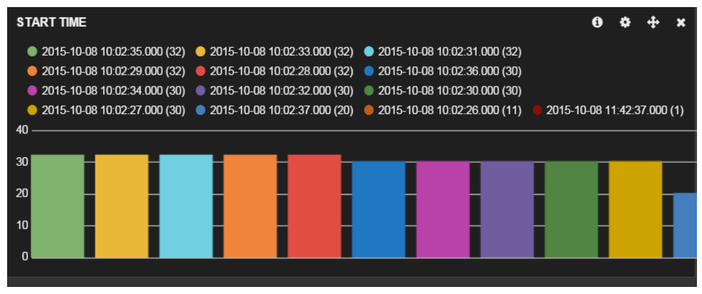

CONCLUSION

Thanks to the introduction of the ELK stack, all the log messages have been unified into a common format and the data storage is centralised. The management and analysis of all these data has greatly improved, users have created their own dashboard according to their needs.
The bug diagnostics has been improved a lot thanks to the ELK stack; all the data logs are centralised in a single application and errors can be identified easily.
The time the developers spent identifying bugs under the RADE framework has been reduced.
One of the improvements we have in mind is to add an access control to avoid interactions of one user's dashboards with other users' dashboards and also to add a security layer on top of all the stored data.
We are also planning to add new graphical components in Kibana and to extend it to be used in other websites outside the ELK stack.

CERN EN ENGINEERING DEPARTMENT ICE INDUSTRIAL CONTROL