# Development of site-oriented Analytics for Grid computing centres

## Introduction

The field of analytics, the process of analysing data to visualise meaningful patterns and trends, has become increasingly important to a wide range of scientific applications. In this study we explored the components necessary to build a distributed Worldwide LHC Computing Grid (WLCG) analytics platform using Grid resources at Glasgow and Edinburgh, two members of the NGI UK distributed Tier-2 Scotgrid. The motivation was in part to use the wider coverage of data to identify common workload trends to be identified and perform error diagnosis not apparent from an individual Grid site perspective. The platform was developed with a focus on providing a site-oriented solution with the intention of complementing ongoing analytics development from larger Virtual Organisations (VOs).

## Approach

The *Elasticsearch*, *Logstash* and *Kibana* (ELK) stack is a combined analytics, logging and visualisation platform that has risen in prominence in the HEP community. Unstructured data from numerous sources can be gathered, curated and then visualised through a single extensible interface. In our approach we combine data collected from ELK stack instances located at the candidate sites into a prototype “regional” ELK instance.

## Data flow

Logstash ingests selected near real-time logging information from site Computing Elements (CREAM CEs) and batch system accounting logs (currently GE, with Condor in progress) by the deployment of *logstash-forwarder* on each service host. Incoming data is parsed and filtered by logstash before being sent to a site Elasticsearch service. A subset of data is forwarded by logstash to the regional ELK instance for further processing and storage (Figures 1 and 2).

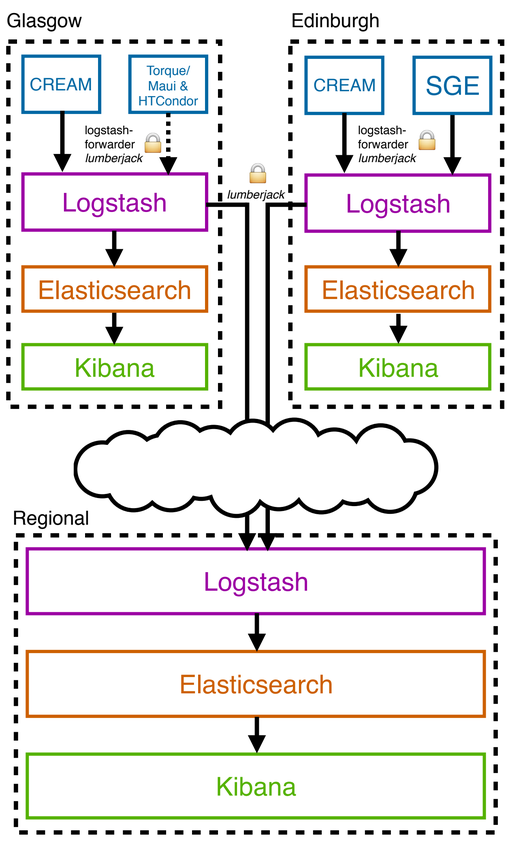


Figure 1 | Overview diagram showing the components of regional ELK model

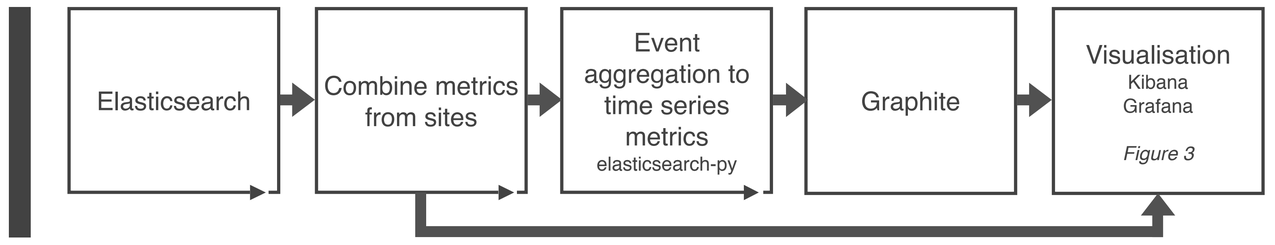


Figure 3 | Process diagram for the derivation of metrics and their aggregation for use in time series plots

Figure 2 | Visualisation of the combination of selected and parsed CE logs from two sites

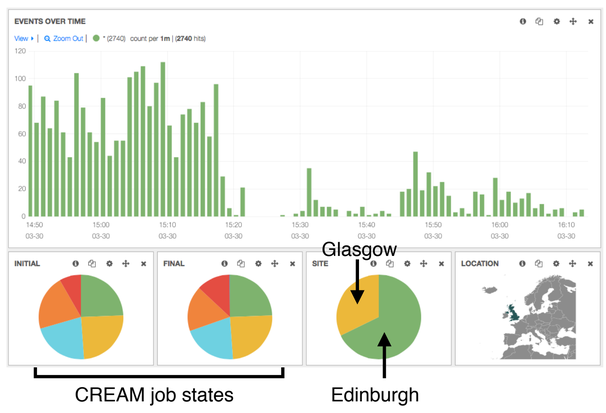


Figure 4 | Rolling average queue times per VO and single-core vs 8-core performance

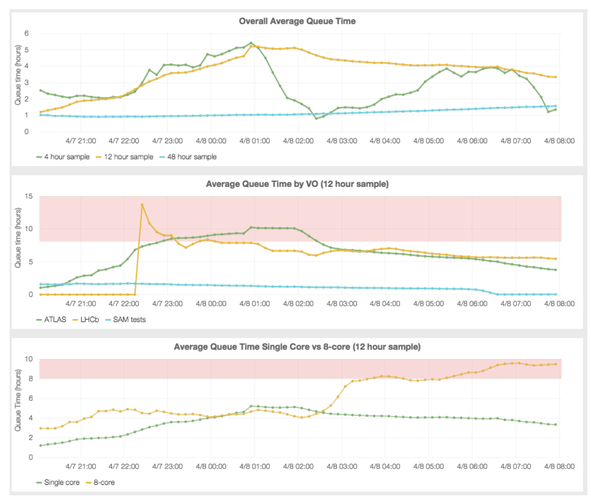


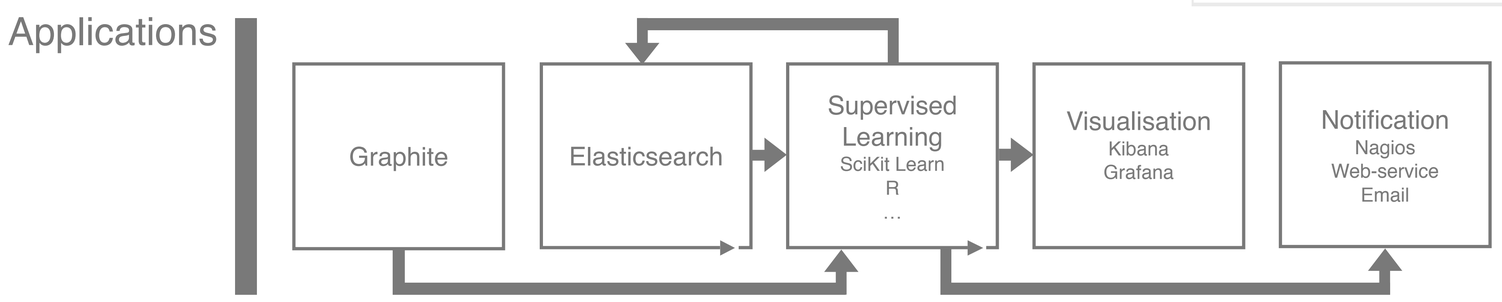


Figure 5 | Process diagram showing analytics applications

## Security

It was important to consider security and privacy measures from the outset:

- Sensitive user data including local account details, delegation IDs and IP addresses were removed prior to transport off the site domain by dedicated logstash filters and by the initial choice of logging data
- Data transmitted between all logstash services were sent through an encrypted channel using the *lumberjack* protocol, using a standard TLS infrastructure with separate self-signed host certificates for each instance
- Access to each Kibana instance was exclusively over HTTPS and restricted to an explicitly listed set of valid grid credentials

## Applications

Once the data of interest was stored in a consistent fashion it was possible to perform more detailed analytics to gain further insight on site operations. Efficient searching and retrieval through the Elasticsearch API and low level clients (e.g. elasticsearch-python) allowed analysis can be continually applied to incoming data (Figure 3). The generated time-series output can be further stored and visualised in auxiliary monitoring tools.
As an example, Grid job metadata stored in the batch system accounting logs were used to extract a rolling average of per-VO job queue times and to inspect the relative scheduling performance of jobs requesting different CPU resources (Figure 4). This was useful for tracking workload throughput performance when providing resources to multiple VOs and (in the case of Edinburgh) sharing underlying resources with non-Grid projects.

## Future Work

The ELK stack has been shown to be effective in developing a distributed site-oriented analytics platform; we will now look towards developing best practices for a production-level solution that could be deployed at other participating Grid computing centres. Further effort will be required in developing robust security and resiliency models, as well as working with the HEP and WLCG communities to expand the range of collected data sources. We have also started to harness supervised machine learning techniques to automatically identify emerging trends in site performance (Figure 5) and we will aim to extend this functionality to provide automatic notification and issue pre-emption services.


Andrew Washbrook[1], David Crooks[2]
Gareth Roy[2], Sam Skipsey[2], Gang Qin[2], David Britton[2]
[1]University of Edinburgh [2]University of Glasgow