Scaling Up the Deployment and Operation of an ELK Technology Stack

Serhiy Boychenko, Pedro Martel, Brad Schofield, CERN, Geneva, Switzerland

WEPHA139

Since its integration into the CERN industrial controls environment, the SCADA Statistics project has become a valuable asset for controls engineers and hardware experts in their daily monitoring and maintenance tasks. The adoption of the tool outside of the Industrial Controls and Safety Systems group scope is currently being evaluated by ALICE, since they have similar requirements for alarms and value changes monitoring in their experiment. The increasing interest in scaling up the SCADA Statistics project with new customers has motivated the review of the infrastructure deployment, configuration management and service maintenance policies. In this paper we present the modifications we have integrated in order to improve its configuration flexibility, maintainability and reliability. With this improved solution we believe we can propose our solution to a wider scope of customers.

Backend Infrastructure

- Value change and alarm statistics are extracted and calculated on a daily basis from WinCC OA Archive database
- WinCC OA application logs are collected from their respective servers and processed in a multi-layer pipeline
- The processed data is then made persisted in the Elasticsearch and displayed in Kibana dashboards

Data Processing Pipeline

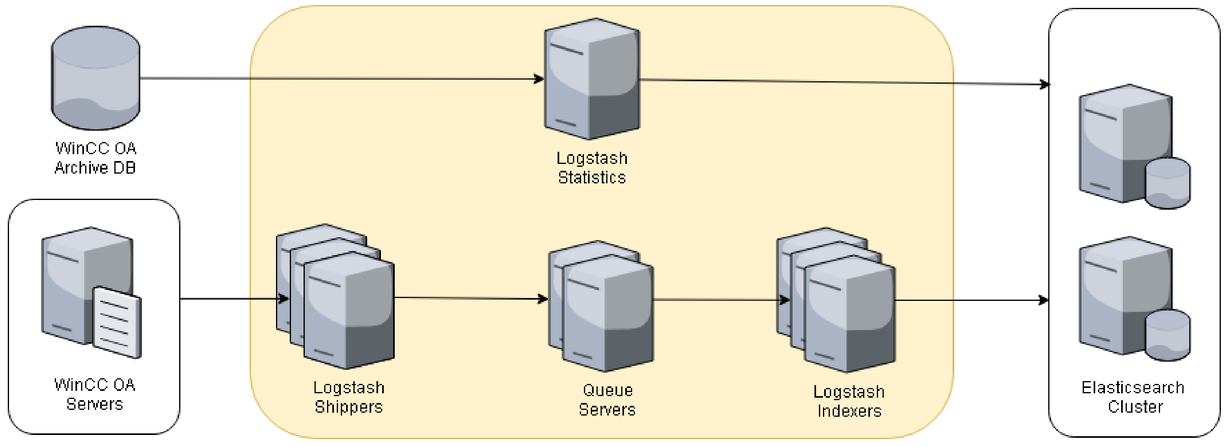

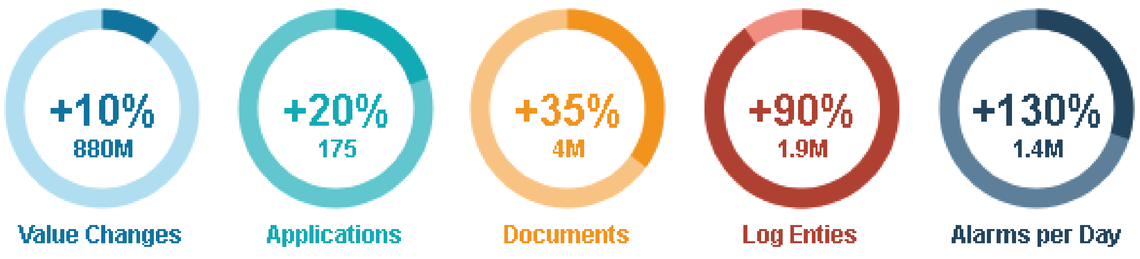

During 2018, the usage of the system was scaled up, introducing more instabilities into the setup.

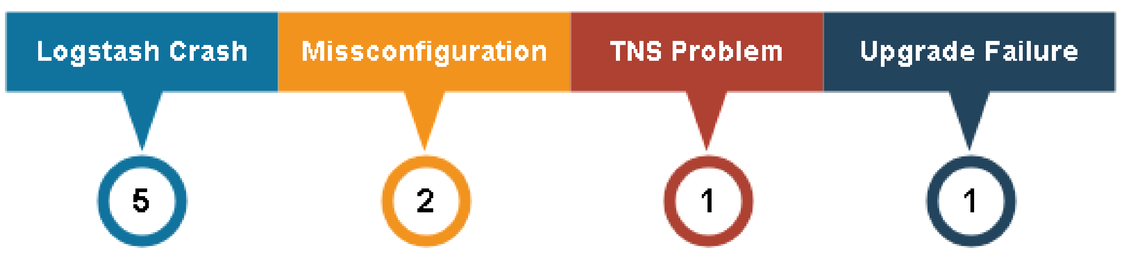

Challenges

- High distribution and heterogeneity of the cluster configuration posed a significant challenge for service availability and reliability
- Continuously increasing data processing volumes along with emerging new use cases required additional effort in order to maintain the service operational and up to date
- Increasing interest for the deployment of similar systems from different WinCC OA user groups at CERN called for improvements in the configuration maintainability and flexibility

Improvement Measures

- Individual Monit (open-source monitoring tool) daemons are deployed on each of the infrastructure machines to monitor the service status, recover from failures and alert when a manual intervention is required
- Different Ansible (open-source provisioning, configuration management and application deployment tool) configurations were developed to deploy the cluster and underlying service functionality
- Multiple Jenkins (open-source Continuous Integration and Delivery Framework) jobs were implemented to orchestrate and automate cluster management process
- GitLab source code repository was used to make produced configurations and scripts persistent

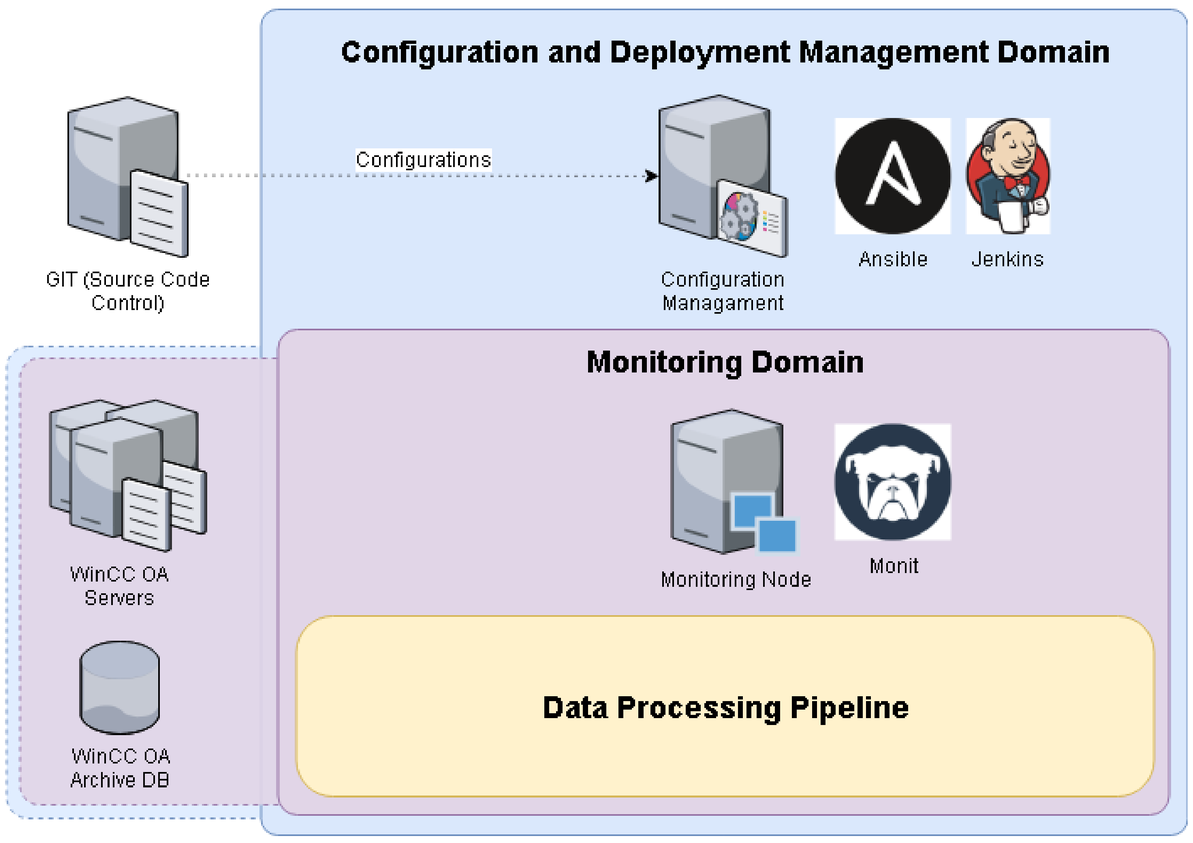

Conclusion

The integration of the monitoring solution allowed us to reduce significantly the time required to detect failures and provided simple yet reliable self-healing mechanisms. Since its integration into the production environment, the number of manual interventions was minimal, in comparison to the previous experience. The improvements of the configuration and deployment management process have proven their efficiency both during the upgrade of the cluster and the setup of a similar infrastructure for external users. We could conclude that the chosen tools integrate well with the Elastic stack, and are viable options for use cases similar to the SCADA Statistics one.

CERN Beams Department
Industrial Controls and Safety Systems Group (ICS)