

## A hybrid system for monitoring and automated recovery at the Glasgow Tier-2 cluster

*Thursday, 20 May 2021 18:00 (30 minutes)*

We have deployed a central monitoring and logging system based on Prometheus, Loki and Grafana that collects, aggregates and displays metrics and logs from the Tier-2 ScotGrid cluster at Glasgow. Bespoke dashboards built on Prometheus metrics give a quick overview of cluster performance and make it easy to identify issues. Logs from all nodes and services are collected to a central Loki server and retained over time. This integrated system provides a full overview of the cluster's health and has become an essential tool for daily maintenance and in the investigation of any issue.

The system includes an automated alerting application that parses metrics and logs and can send notifications when specified conditions are met, and as a further step toward automation, can also perform simple recovery actions based on well known issues and their encoded solutions. The general purpose is to create a more resilient Tier-2 cluster where human intervention is kept to a minimum. Given the funding constraints experienced by many academic research institutions, this promises to free staff from routine tasks, allowing them to address their expertise to more interesting problems.

In this paper, we describe the tools and set-up of the existing monitoring system, the automated recovery methods implemented so far, and the plan for further automation.

**Primary author:** SIMILI, Emanuele (University of Glasgow)

**Co-authors:** ROY, Gareth (University of Glasgow); STEWART, Gordon (University of Glasgow); SKIPSEY, Samuel Cadellin; BRITTON, David

**Presenter:** SIMILI, Emanuele (University of Glasgow)

**Session Classification:** Thurs PM Plenaries

**Track Classification:** Distributed Computing, Data Management and Facilities