

Identifying Opportunities to Improve Efficiency in HPC Clusters

Jordi Blasco
NeSI
Auckland, New Zealand
Email: jordi.blasco@nesi.org.nz

Abstract—Jordi Blasco has developed a new open source monitoring tool which allows the HPC user support teams to identify new opportunities to improve the efficiency of the codes being executed on HPC resources. Earlier adopters of this new tool, and through the continuous monitoring of jobs efficiency, have been able to improve scalability and performance of several codes and workflows. This in turn has accelerated research in those HPC facilities adopting this technology.

In addition, the improvement in the global efficiency of the system has had an effect on overall resource allocation, allowing HPC users to run more jobs and/or to deal with bigger problems.

Traditional tools like Ganglia are not normally capable of representing the metrics required to identify inefficient jobs. Nor are they capable of correlating events in order to identify global issues affecting all jobs in the system. The new monitoring system allows end user support team to evaluate the performance impact on running jobs due to systemic issues such as a high load on the cluster file system or a high rate of hardware errors in the fabric.

Large number of events to analyze requires use of Big Data technologies. Data is gathered using custom codes and aggregated into Elasticsearch, an open source search and analytics engine which has high reliability and proven scalability. Finally, the data is represented through Grafana, which is a leading tool for querying and visualizing large datasets and metrics.

The talk highlights the most relevant cases where NeSI applied the tool to identify tuning opportunities and accelerate research by improving code efficiency.

Index Terms—Performance Analysis, Efficiency, Scalability, Job profiling.

