

Title: Practical implementation of monitoring on Cray system

Abstract:

Monitoring has become an area of significant interest for many Cray sites and a System Monitoring Working Group was established by Cray to support collaboration between sites in monitoring efforts. An important topic within this group is the practical implementation of monitoring. Different sites have implemented monitoring of different aspects of Cray systems, and the aim of this BoF is to share practical implementation recommendations so we can learn from each other's experiences rather than independently repeating them via trial-and-error. This BoF will consist of presentations from sites on lessons learned in practical experience, interspersed with discussion about problems sites are facing and known or proposed solutions to these.

Supporting BoF material:

This BoF targets practical implementation recommendations and discussion and would be complementary to a BoF being proposed on "Automated Analysis and Proactive Response". They could be held back-to-back in a single session.

We have two proposals so far for presentations within the BoF and anticipate gathering more in the coming weeks

"Experiences with Xalt at on Piz Daint at CSCS"

Jean-Guillaume Piccinali, CSCS

We will describe our experience in using Xalt on a Cray XC40/XC50 hybrid supercomputer at CSCS (Piz Daint). Xalt intercepts the user link (ld) and job launcher (srun) lines, maps program name and libraries to modulefile names, records complete list of environment variables, stores the results and provides reporting scripts. Thus, it allows to know which users are running what applications, and how. Based on lessons learned from deploying Xalt at CSCS, we will provide recommendations about the practical setup of this tool on a modern HPC platform and compare the advantages and disadvantages of a relational database such as MySQL. We will give an overview of what we think is a better solution for storing this information in the context of CSCS monitoring infrastructure. We will conclude with a discussion about the different ways that we report this data, and identify areas of improvement with the audience.

"A reference architecture for collecting monitoring data from Cray systems"

Stephen Leak et al, NERSC

NERSC has a sophisticated infrastructure, based around Elasticsearch and RabbitMQ, for collecting and visualizing large volumes of data collected by different monitoring tools. Few sites have the resources to replicate such a project, so the CUG System Monitoring Working Group has assembled a reference architecture built from Docker containers and modeled after the NERSC infrastructure. This allows a site to build a basic monitoring system on one or a few spare servers and customise and scale it according to local resources and priorities. We will present an overview of the architecture along with templates for its components and practical tips for setting up and using it.