LA-UR-20-25796
Approved for public release; distribution is unlimited.

Title: A Virtual Cluster Monitoring Toolkit for Bottleneck Analysis

Author(s): Frumkin, Natalia
Marquardt, Christian Michael

Intended for: HPC Intern Showcase, 2020-08-13 (Los Alamos, New Mexico, United States)

Issued: 2020-07-31

# A Virtual Cluster Monitoring Toolkit for Bottleneck Analysis • Natasha Frumkin & Christian M. Marquardt August 3rd, 2020

Monitoring computing clusters and detecting potential bottlenecks is a necessary part in designing and engineering efficient, reliable systems. HPC-Collab, a tool for constructing virtualized HPC clusters, automatically configures a virtual cluster with a specified cluster topology and software configurations. However, HPC-Collab suffers from scalability issues. Due to long installation times of about two hours for a fifteen node cluster, using HPC-Collab for spinning up proof-of-concept experimental clusters is time-intensive.

Our project collects useful metrics from the virtual cluster nodes as they are provisioned, and correlates virtual cluster activity with the host. We designed a monitoring infrastructure which collects relevant system data from the virtual clusters as they are created. At the same time, we collect host machine metrics to correlate activity from virtualized nodes with underlying IO patterns, network traffic, and memory usage. Using additional helper scripts, data from multiple virtual nodes is gathered and visualized through custom dashboards. This enables system performance analysis at each stage of the process.

From our experiments, we have identified major bottlenecks are primarily due to network bandwidth and large overhead from the underlying virtualization providers. Contrary to originally inferred, we have proven that neither random IO patterns of guest-to-host layered file systems nor CPU usage are bottlenecks in the provisioning process. Additionally, we have provided quantitative measurements to verify how long each provisioning takes as well as where in the provisioning process we see large fluctuations in network traffic and RAM usage. As a result of our work, cluster provisioning time was reduced by 50% for the standard reference cluster model.

In the future, we plan to provide fine-grained graphs which we directly match up with the phases in virtual cluster provisioning that are particularly slow. Ultimately, we hope to integrate automatic monitoring into the HPC-Collab project and provide developers with a visualization tool for quick and intuitive bottleneck analysis. This aligns with hpc-collab's project goal: promoting engineered, rather than artisanal, cluster construction.