NEXTGenIO Workshop on applications of NVRAM storage to exascale I/O



Contribution ID: 30

Type: Oral presentation

A top-down performance analysis methodology for workflows

Wednesday, 25 September 2019 16:45 (30 minutes)

Scientific workflows are well established in parallel computing. A workflow represents a conceptual description of work items and their dependencies. Researchers can use workflows to abstract away implementation details or resources to focus on the high-level behaviour of their work items. Due to these abstractions and the complexity of scientific workflows, finding performance bottlenecks along with their root causes can quickly become involved. This work presents a top-down methodology for performance analysis of workflows to support users in this challenging task.

Our work provides summarized performance metrics covering different workflow perspectives, from general overview to individual jobs and their job steps. These summaries allow to identify inefficiencies and determine the responsible job steps.

In addition, we record detailed performance data about job steps, enabling a fine-grained analysis of the associated execution to exactly pinpoint performance issues. The introduced methodology provides a powerful tool for comprehensive performance analysis of complex workflows.

Keywords

Presenter: WILLIAMS, Bill (GWT-TUD GmbH)

Track Classification: NEXTGenIO Workshop on applications of NVRAM storage to exascale I/O