HPC22

A Holistic Approach to Solving Petascale Problems

C. Fazzino* (ExxonMobil)

SUMMARY

As PetaScale high performance computing systems are deployed in the energy sector the full capability of these systems is often not realized. These systems are normally rated by their compute capacity in terms of Rmax , peak FLOPS. This measure says little about how well a system solves a particular problem. Given Supercomputers high cost and the strategic importance of the business decisions they influence, the consequences of delivering an underutilized or over configured system is detrimental to the investment's value. For Supercomputers to deliver their full value, a comprehensive approach is recommended that considers the problem domain, human factors, software engineering, and system architecture. This paper will discuss how interplay of these factors can work to maximize the value delivered for the supercomputer investment.

Introduction

As PetaScale high performance computing systems are deployed in the energy sector the full capability of these systems is often not realized. These systems are normally rated by their compute capacity in terms of Rmax, peak FLOPS. This measure says little about how well a system solves a particular problem. Given Supercomputers high cost and the strategic importance of the business decisions they influence, the consequences of delivering an underutilized or over configured system is detrimental to the investment's value. For Supercomputers to deliver their full value, a comprehensive approach is recommended that considers the problem domain, human factors, software engineering, and system architecture. This paper will discuss how interplay of these factors can work to maximize the value delivered for the supercomputer investment.

Method and/or Theory

High performance computing problems vary greatly in terms of data size, computational requirements, inter-process communication, and workload. These characteristics have a significant impact on the overall performance of the system. A slow interconnect on a system which requires high volumes of inter-process communication will leave node resources idle. Accelerators are not effective when internal memory sizes are too small to hold the data needed to complete an operation. Therefore, it is important to have a thorough understanding of these requirements when designing and operating a high performance computing system.

System architecture and management will positively or negatively impact the capability of a machine. An architecture that balances the system resources to the characteristics of the problem domain will consistently perform better on those problems than a system with higher LINPACK numbers. During the design phase it is wise to bench mark the performance of potential systems with your applications and data to ensure the system will perform as expected. Once in operation it pays to monitor application use and system utilization and seek user feedback to ensure the system continues to deliver its' maximize value. System management tools are needed in HPC to help maintain the overall health of the system by automating housekeeping and helping enforce good user behavior. For example, providing job queues with different configurations can be established to ensure jobs behave well and that the system is fully used. Tuning the system architecture and management to the software architecture and business workflow increases the resource utilization and reduces waste.

An in-depth understanding of the software engineering stack and how it is integrated within the system is crucial to attaining scale. The overriding objective for many HPC applications is getting working code into production. Little effort is normally given to usability, user productivity, performance optimization, integration, reuse, project and job management, and scalability. Sometimes over-engineering proprietary applications limits the ability for the system management tools to perform their tasks. Introducing features into software that assists users in optimizing resource utilization can enable the average user to intelligently scale problems to the available resources. Each one of these factors limits the ability to scale applications in production. However, time spent in this area will yield step change improvements to job scalability and through-put.

Users' skill levels and their motivation to work in the general interest also impact scalability. In this environment any user can affect everyone's performance and diagnosing problems at this scale can be difficult. Skilled users are less likely to cause problems and are more productive because they are able to: effectively decompose their problem; monitor job status; quickly identify issues and correct them; and full leverage system resources. Time spent training and helping users apply these skills takes time, but is the most effective way to raise organizational capability. You also need users to be motivated to not waste system resources and continually focus on achieving marginal cycle time improvements from their applications. System governance is important to ensure the system

resources are being used for the most important work and to establish policies that encourage users to work in the general interest.

Conclusions

High performance computing systems are costly to operate and often times do not perform as expected. Comparing systems based on peak FLOPS doesn't necessarily tell much about what value a system will deliver. A holistic approach that aligns the system architecture, software architecture and human factors with the problem domain is essential to maximize the value from your HPC investment.