Application to attend NSF PI Meeting on "The Science of Cloud Computing"

Applicant: Vaidy Sunderam
Emory University
vss@emory.edu

Topics: 8) Cloud Self-Monitoring and Autonomic Control
9) Cloud Interoperability and Standardization

Current Research Related to Cloud Computing:

Our current work on Clouds is the latest in a long-running series of research projects on various aspects of heterogeneous metacomputing, resource aggregation and sharing, and autonomic grid systems, most recently supported by NSF grants ACI-0220183 and CNS-0720761. In early foundational work (PVM, EcliPSe, and PIOUS), our efforts focused on systems for network parallel processing, message passing interfaces, concurrent I/O frameworks, loosely coupled multithreaded concurrency, and hybrid client-server/asynchronous environments. Other projects including Harness and IceT focus on collaborative and component based metacomputing. Our current research efforts are the Unibus system and the H2O project. These projects address important issues in grids, clouds and metacomputing systems arising from (1) heterogeneity in resource types, (2) difficulties in preparing and staging applications, and (3) intrinsically high levels of dynamicity common in diverse computing resources. Strategies to reconcile different resource classes, and to dynamically provision resources that become available during application execution are the focus of the current ADAPT project continuing our recent work on Unibus. Cloud systems present strong motivation for addressing these goals. Given the variability in Cloud platforms, usability is a challenge, particularly for legacy science and engineering applications. Multiple access techniques for IaaS clouds present considerable logistical hurdles for application deployment thereby raising barriers to flexible switching between the most appropriate resources for different runs. PaaS clouds are of course closely coupled to a specific programming model, but translations were possible, greater effectiveness might ensue. The ADAPT project aims to address these issues by dynamically creating or assembling mediators perform reconciliations between applications and target Cloud platforms as well as grids and on-premise resources.

Application to attend NSF PI Meeting on "The Science of Cloud Computing"

Applicant: Vaidy Sunderam
Emory University
vss@emory.edu

Future Research Problems:

Cloud technologies are emerging as a promising approach to *computing as a utility*, a vision that has great appeal in various domains. However, the *usability* of cloud platforms for efficiently executing science and engineering applications is proving to be a formidable challenge. Primary among the obstacles to widespread adoption of cloud computing in S & E is the potential mismatch between programming models in which (legacy and current) applications are written, and the facilities provided by target cloud platforms. We intend to investigate novel approaches to enhance executability of applications on varied computational back-ends including different types of clouds as well as traditional (local) resources. We will propose abstractions, to be realized as lightweight middleware, for *dynamic adaptation* of computing models to specific resources, that will: (1) permit on-premises applications to seamlessly execute on multiple platforms including clouds; and (2) enhance the cross-execution ability of cloud applications.

A second area of great interest is the issue of *interconnection* within and across Clouds. Currently, cloud providers focus on computational capability and storage capacity. However for most S & E applications, *connectivity* and *interaction* are crucial. To execute a parallel application on an IaaS cloud, for example, allocation of computation and storage resources is straightforward but only general purpose networking, that too low-capacity, is available for message passing. We believe that Cloud providers will eventually supply interconnect capability as a service but even then, applications will have to evolve new mechanisms for dealing with the important issue of inter- and intra-Cloud program interaction. We propose to investigate new paradigms for *expressing, implementing, and using connectivity facilities* in current and emerging Cloud platforms.