

Model-Based Services for Computing Grid Performance Analyses & Tuning

Rob Simmonds and Brian Unger
Grid Research Centre
University of Calgary, Canada
E-mail: simmonds@cpsc.ucalgary.ca

Abstract

Grid computing environments typically federate heterogeneous resource “clusters” belonging to several organizations. To fully realize the promise of a grid environment, it is necessary to support analysis tools that help obtain insights into the behavior of individual clusters. The Grid Research Centre at the University of Calgary (<http://grid.ucalgary.ca/>) is developing a cluster service called “replay” that simplifies the development and maintenance of such tools. The service provides a single, consistent, model-based interface for obtaining current and historical information about a cluster.

Replay will manage multiple views of a cluster which enables tools to obtain data from an existing cluster as well as information that shows how a cluster might have behaved under alternate “simulated” configurations and workloads. Using information collected at a series of time instants, tools can “replay” the activity observed over a given time period. This replay can reproduce actual historical data or observe simulated alternatives. The system will support “rollback” to recover the state of a cluster at a particular time.

The model-based nature of the service also permits tools to obtain information transparently regardless of the underlying cluster. Another unique feature of replay is its use of different mechanisms to manage information that typically changes infrequently and information that can change in a more dynamic, continuous manner.

This work is primarily motivated by the analysis tools developed for WestGrid, a grid-enabled high performance computing facility in Western Canada.