

LQCD Workflow Execution Framework: Models, Provenance, and Fault-Tolerance

Monday 23 March 2009 08:00 (20 minutes)

Large computing clusters used for scientific processing suffer from systemic failures when operated over long continuous periods for executing workflows. Diagnosing job problems and faults leading to eventual failures in this complex environment is difficult, specifically when the success of whole workflow might be affected by a single job failure.

In this paper, we introduce a model-based, hierarchical, reliable execution framework that encompass workflow specification, data provenance, execution tracking and online monitoring of each workflow task, also referred to as participants. The sequence of participants is described in an abstract parameterized view, which is translated into a concrete data dependency based sequence of participants with defined arguments.

As participants belonging to a workflow are mapped onto machines and executed, periodic and on-demand monitoring of vital health parameters on allocated nodes is enabled according to pre-specified rules. These rules specify conditions that must be true pre-execution, during execution and post-execution.

Monitoring information for each participant is propagated upwards through the reflex and healing architecture, which consist of hierarchical network of decentralized fault management entities, called reflex engines. They are instantiated as state machines or timed automata that change state and initiate reflexive mitigation action(s) upon occurrence of certain faults.

We describe how this cluster reliability framework is combined with the workflow execution framework using formal rules and actions specified within a structure of first order predicate logic that enables a dynamic management design that reduces manual administrative workload, and increases cluster-productivity. Preliminary results on a virtual setup with injection failures are shown.

Presentation type (oral | poster)

oral

Authors: DUBEY, Abhishek (Vanderbilt University); KOWALKOWSKY, James (Fermilab); SIMONE, James (Fermilab); PICCOLI, Luciano (Fermilab)

Presenter: PICCOLI, Luciano (Fermilab)

Session Classification: Poster session

Track Classification: Distributed Processing and Analysis