



Contribution ID: 417

Type: Poster

Event Workflow Management System - A SaaS Solution for Massively Divisible and Distributed Workflows

How does one take a workload, consisting of millions or billions of tasks, and group it into tens of thousands of jobs? Partitioning the workload into a workflow of long-running jobs minimizes the use of scheduler resources; however, smaller, more fine-grained jobs allow more efficient use of computing resources. When the runtime of a task averages a minute or less, severe scaling challenges due to scheduling overhead can surface. Employing jobs that run for several hours, each with a large input file comprising a bundle of tasks, is effective in ideal situations. However, given the heterogeneity of available distributed resources and limited control of task-job matching, runtimes can vary widely.

The Event Workflow Management System (EWMS) augments HTCondor to solve this issue. EWMS implements a pilot-based paradigm where each worker, running inside an HTCondor execution point, connects to a message broker and executes many individual fine-grained tasks. This adaptive design increases task throughput while incorporating additional fail-safe features. In addition, EWMS manages workflow scheduling, enables real-time worker scaling, and exports a public-facing interface for user accessibility. Here, we outline the EWMS technique, detail science driver workflows from the IceCube experiment, and provide system usage metrics.

Primary author: EVANS, Ric (Wisconsin IceCube Particle Astrophysics Center)

Co-authors: RIEDEL, Benedikt (University of Wisconsin-Madison); AYDEMIR, Brian (Morgridge Institute for Research); BOCKELMAN, Brian Paul (University of Wisconsin Madison (US)); SCHULTZ, David (University of Wisconsin-Madison); LIVNY, MIRON

Presenter: EVANS, Ric (Wisconsin IceCube Particle Astrophysics Center)

Session Classification: Poster session

Track Classification: Track 4 - Distributed Computing