

Supporting Phase Dataset Intervals for Job Recovery:

Proposed Design

By

Andrew Baranovski, Gabriele Garzoglio, Parag Mhashilkar

Abstract

Submitting recovery jobs for Monte Carlo production is not very user friendly. The user has to go through a series of manual steps before a recovery job is submitted. This document describes the problem in detail and proposes solution to ease the recovery job submission and tracking process.

Introduction

For Monte Carlo production user can submit two types of jobs -

1. Production jobs where the events are generated on the site.
2. Production jobs which get the events from the phase dataset file.

In this document we describe the recovery of jobs in the later case. Following sections describe how we intend to improve the efficiency of the recovery process by using the information available in SAM and adding extra functionality in SAMGrid.

Problem Description

In order to submit the recovery job user has to browse JIM Monitoring page for the failed grid job and find the list of failed local jobs and the corresponding events. After that the user has to submit a separate recovery job for processing range of adjacent events.

We have identified two drawbacks in the above operational model -

1. There is lack of sufficient support in operational tools to acquire the range of events by querying SAM. This leads to considerable manual work on part of the user.
2. Even if these event ranges are available, there is no support in SAMGrid to accept them as input for the Monte Carlo recovery job. One can specify the start even and offset but this is barely usable to automate the process and make it user friendly.

For example, consider an input file with 2500 events (Figure 1) split into blocks of 250 events. Grey blocks represent the event ranges that need to be reprocessed due to the failure in the production job. Today, in order to submit a recovery job user has to submit three grid jobs for the event ranges 1-250, 501-1000 and 1751-2000

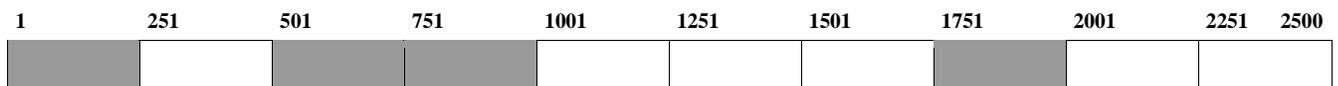


Figure - 1

Proposed Solution

The manual steps discussed in the problem description above can be automated by extending support in the operational tools coupled with support in the SAMGrid infrastructure. This will also make tracking of recovery jobs more user friendly by combining the individual recovery job into one grid job.

Proposed changes to operational tools

Operational tools could be modernized to take the request id of the Monte Carlo run over which user will run the recovery job. Using this request id tools can query the SAM database to get *phase_dataset_intervals*. The parameter *phase_dataset_intervals* is the event ranges that need to be processed in the recovery job.

You can get the *phase_dataset_intervals* using following algorithm –

```
# JIMJobId = Global JobId of the samgrid Job
# StartEvent, EndEvent = Event ranges that are being processed by the job

get_phase_dataset_intervals (JIMJobId, StartEvent, EndEvent)
    requested_range = set(StartEvent, EndEvent)
    processed_events = set of processed events for the JIMJobId obtained from SAM
    unprocessed_events = requested_range - processed_events
    phase_dataset_intervals = convertToIntervals(unprocessed_events)
    return phase_dataset_intervals
# end get_phase_dataset_intervals
```

Thus, in case of figure – 1

phase_dataset_intervals = 1-250, 501-1000, 1751-2000

Proposed changes to SAMGrid

SAMGrid needs to be modified to support *phase_dataset_intervals*. Once the intervals are available to the SAMGrid system, it will submit local jobs that will process the specified intervals for recovery.

Figure – 2 shows the recovery job submitted for figure – 1 with 200 events/job.

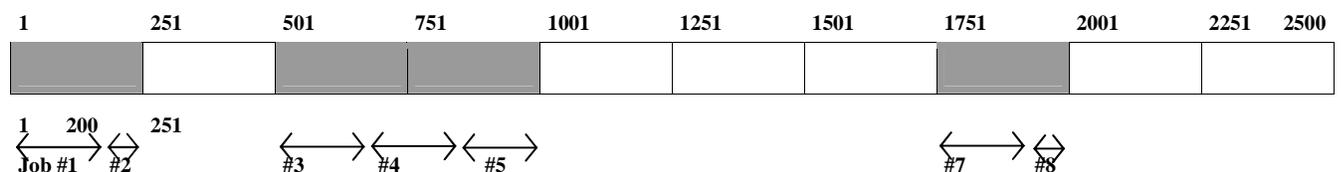


Figure - 2