



POMS Introduction - Dune

Marc Mengel, Stephen White

POMS

August 2019



What is POMS?



- Provides a simple and transparent interface between users and the complexity of the grid.
- A service to assist production and analysis of experiments in their MC production and DATA processing.
 - As the quantity of data originated by the running experiments greatly increases, the ability of simplifying the steps in data processing and management has become more and more appealing to the users
- A web service interface, enabling automated jobs submission on distributed resources according to customers' requests and subsequent monitoring and recovery of failed submissions, debugging and record keeping.
 - The ultimate goal is the most efficient utilization of all computing resources available to experiments, while providing a simple and transparent interface between users and the complexity of the grid.

What does POMS do?



Job Launches: allowed

Stephen White

dune production

Obtaining a POMS Account

POMS Access is supplied through FNAL services accounts. The account must be registered in FERRY as analysis, production or superuser. If an account does not have access or needs a different level of access, follow the steps in the help link below.



Creating a Campaign

We recommend creating using the campaign editor. A new campaign can be created from the Campaigns page with the add button or cloning and modifying exiting ones. For an understanding on how to do this please read the user documentation in the provided help link.



Monitoring

POMS provides monitoring through Landscape and internal graphs. All job related information is available in FifeMon. Internal charts exists to track campaigns. These are available when viewing a specific stage.



Effcency


Debugging

Log files are available for debugging once a stage is launched. They are accessible from the 'Campaign Stages' by viewing the details of that stage.



POMS Statistical Overview

Jobs By Group



Overview

Fife_Utils

POMS utilizes fife_launch and fife_wrap from the Fife_Utils tool set. This provides a set of commands and interfaces with JOBSUB passing the parameters needed to run the jobs. Details of this can be found in the help link below.



FifeMon

FifeMon gathers information from many different sources and graphs it on a common timeline to help experiments understand their computing usage and identify problems.

Overview

Shifter

SAM

SAM is a data handling system organized as a set of servers which work together to store and retrieve files and associated metadata, including a complete record of the processing which has used the files.

Status

Wiki

Jobsub

Jobsub is a suite of tools to manage batch/grid submission. These tools are designed to simplify the job submission process by defining interfaces for experiments, integrating grid tools, protecting shared resources.

Batch Detail

Wiki

dCache

dCache is disk caching software capable of massive data throughput. It stores and retrieves large amounts of data, distributed among a large number of heterogeneous disk server nodes.

Overview

Wiki

- Submits grid jobs
- Tracks job submissions through Landscape
- Organizes job submissions into Campaigns with Stages
- Has a Graphical Editor for Campaigns and their Stages.
- Records how and when submissions were run
- Assists in analyzing job failures with plots, charts, and easy access to log files.
- Can also be used for Analysis, providing a common architecture for both production and analysis
- Experiment Superusers see all computing in one place

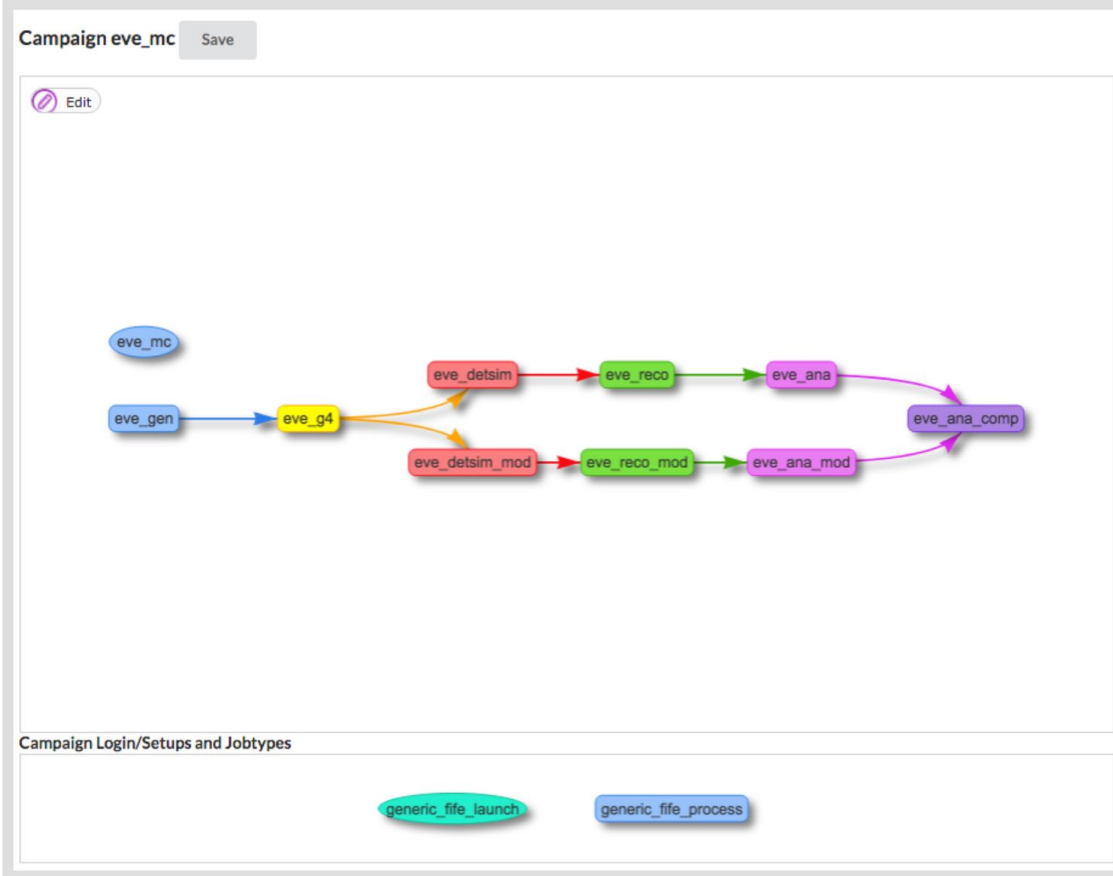
POMS and SAM

- POMS is able to leverage external mgmt systems, which is currently done using SAM. Thus it works best when:
 - Input (if any) is a SAM dataset, fetched from SAM Project
 - OR output files added to a “poms_...” dataset
 - Output files are declared to SAM (possibly by F-FTS)
 - Output files are descendants of inputs in SAM
- POMS Can be used without SAM
 - Submission scripts must find files on their own
 - Recovery launches must find failed files on their own
 - Dependency launches cannot wait for availability
 - File based reports don't work
- Sam code factored into one module, (plus split types)

POMS and Future Data Management

- POMS will be able to use future data management with little (or no) change
- Rucio for
 - Data movement
 - Location tracking
- Future Data management will provide
 - Metadata information
 - Logical (metadata query) datasets
 - project/station functionality

POMS Campaigns



POMS organizes data processing into campaigns with stages, where each stage performs a specific step of processing.

Example of a campaign, shown in the GUI Editor, with stages connected by their dependencies.

Login and job configurations.

POMS Campaign Details

Details of each component are editable via the GUI editor.

The screenshot displays the POMS Campaign GUI editor. At the top, a workflow diagram shows a sequence of stages: `gen` → `g4` → `detsim` → `reco1` → `reco2` → `mergeana`. A blue oval highlights the campaign name `vito_prod_muminus_0-2.0GeV_isotropic_uoone` above the `gen` stage.

Below the workflow, there are two configuration panels:

- dependency gen -> g4**: A panel with a text field for `file_pattern_1` containing `%%`, and buttons for `Reset`, `OK`, and `Cancel`.
- Campaign Login/Setups and Jobtypes**: A panel with a blue oval highlighting the login setup `vito_login`.

The main configuration panel, titled **campaign_stage**, is open for the `g4` stage. It contains the following fields:

name	g4
vo_role	Analysis
software_version	v07_07_00
dataset_or_split_data	None
cs_split_type	None
completion_type	complete (default)
completion_pct	95
param_overrides	[["-stage ", "g4"], ["-Oglobal.sam_dataset=", "%
test_param_overrides	[["-stage ", "g4"], ["-Oglobal.sam_dataset=", "%
login_setup	vito_login (default)
job_type	vito_prod_muminus_0-2.0GeV_isotropic_u
merge_overrides	False
stage_type	test

Buttons for `Reset`, `OK`, and `Cancel` are located at the bottom of the **campaign_stage** panel. `Edit` buttons are also present next to the `cs_split_type`, `param_overrides`, and `test_param_overrides` fields.

POMS Campaign Management

All campaigns, production & analysis, are viewed and managed from the same page.

Active
 Inactive

Mine
 Others

Analysis
 Production

[Update View](#)

[Mark Campaign Active/Inactive](#) [Tag/untag Campaigns](#) [Add](#) [Pick .ini file](#) and [Upload](#)

Select All <input type="checkbox"/>	Campaign Name	id	Active	Submissions Running	Submission History	Launch	Dependencies	GUI Editor	Clone Campaign	.ini File	Tags	Rename Campaign	Delete Campaign	Creator
<input type="checkbox"/>	hakan_RITM0848148_0	3572	Yes	0										hsolak
<input type="checkbox"/>	patch_metadata_hv_values	3245	Yes	1										dingpf
<input type="checkbox"/>	protodune-sp_keepup_cosmics_v07_08_00_05	3224	Yes	0										kherner
<input type="checkbox"/>	protodune-sp_keepup_cosmics_v07_08_00_05_7.5ms	3261	Yes	0										kherner


- ★ Actions on campaigns depend on the users privileges.
- ★ Superusers can start, stop and edit all campaigns.

POMS Campaign/Stage Details

Campaigns can be drilled down to stages, submissions and even specific log files.

Campaign Stage `bnb_reco1_B` 3904
in Campaign `data_bnb_reco1_mcc9` 3216

Batch Job Status



status	total	percentage
Cancelled	0	0%
Fail	981	57%
Held	0	0%
Success	736	43%

Active Job Status

idle
running
max idle

Batch Job Efficiency

Submission #472724
jobsub.job.id: 21249982@jobsub01.fnal.gov
Campaign Stage `bnb_reco1_B` 3904
in Campaign `data_bnb_reco1_mcc9` 3216

Information

Submission was a regular launch
[View in Landscape](#)
[View the SAM Project](#)
[View the Job Logs](#)

Actions

[Hold Condor Jobs](#)
[Release Condor Jobs](#)
[Kill Condor Jobs](#)
[Kick to Located](#)
[Re-Launch Submission on dataset uboonepro_prod_bnb_swizzle_optfilter_open_inclusive](#)
[Launch Recovery for Submission](#) [consumed_status](#) [Launch](#)

History

2019-07-30 11:20:44.496124-05:00 Idle
2019-07-30 11:31:14.621813-05:00 Running
2019-07-30 11:48:33.742895-05:00 Completed
2019-08-01 11:51:20.901963-05:00 Located

Submission Details

submission.id: 472724
jobsub.job.id: 21249982@jobsub01.fnal.gov
campaign.stage.id: 3904
creator: 438
created: 2019-07-30 11:13:01.591758-05:00
updated: 2019-08-01 11:51:20.901963-05:00
command_executed: jobsub_submit --group=uboone --role=Production file:///tmp/tmp50E2f8/submit.dag uboonepro_prod_bnb_swizzle_optfilter_open_inclusive_v3_B_recur_20190730111353
project: login_setup_snapshot_id: 40702
campaign_stage_snapshot_id: 86557
job_type_snapshot_id: 42452
recovery_position: None
recovery_tasks_parent: None

Campaign Stage Snapshot

Results of Output of submit command

```
+ export KRB5CCNAME=/tmp/krb5cc_poms_submit_uboone
+ KRB5CCNAME=/tmp/krb5cc_poms_submit_uboone
++ klist -kt /home/poms/private/keytabs/poms.keytab
++ tail -1
++ sed -e 's/.*/'
+ kinit -kt /home/poms/private/keytabs/poms.keytab
poms/cd/pomsgpvm01.fnal.gov@FNAL.GOV
+ ssh -tx uboonepro@uboonegpvm02.fnal.gov
Pseudo-terminal will not be allocated because stdin is not a
terminal.
Setting up old larsoft UPS area...
/cvmfs/fermilab.opensciencegrid.org/products/larsoft/
Setting up new larsoft UPS area...
/cvmfs/larsoft.opensciencegrid.org/products/
Setting up uboone UPS area...
/cvmfs/uboone.opensciencegrid.org/products/
Active ups products:
art          v3_00_00      -f Linux64bit+2.6-2.12 -q
e17:prof    -z /cvmfs/larsoft.opensciencegrid.org/products
artdaq_core v3_04_01      -f Linux64bit+2.6-2.12 -q
```

statistics for last 30 days

Reports/Status

- [Production Shifter Page \(Landscape\)](#)
- [POMS Campaign Stage \(Landscape\)](#)
- [Campaign Stage Stats \(Landscape\)](#)
- [Submissions \(Landscape\)](#)
- [Campaign Stage Submissions \(1wk\)](#)
- [Campaign Stage Submission Files](#)




Actions

- [Edit Campaign](#)
- [Launch Campaign](#)
- [Launch Campaign](#)
- [Kill Running](#)
- [Hold Running](#)
- [Release Held](#)
- [Schedule Future](#)
- [Queue Future](#)

Campaign Stage

Name: `bnb_reco1_B`
Id: 3904
Experiment: uboone
Dataset: dataset
Software Version: `v08_00_00_01a`
VO Role: Production
Param Overrides: [
[-xml, /uboone/app/users/uboonepro/poms/xml/MCC9/December_2018/data_bnb_reco1_mcc9-v01a.xml]
[-stage, bnb_B]
Split Type: draining
Last Split: None Reset
Created: 2019-07-23 14:46:18.565130-05:00
Creator: renlu23
Updated: None
Updater:

POMS Submission Details

Submission Id	Jobsub Id	Commands	Status (History/Current)	Submitted	Submitter
432572	14721466@jobsub01.fnal.gov	▶ 🗑️ ▶▶	 Located	2019-01-07 14:09:57.542480-06:00	renlu23
432563	14773137@jobsub02.fnal.gov	▶ 🗑️ ▶▶	 Located	2019-01-07 13:19:33.947228-06:00	renlu23
432556	14772054@jobsub02.fnal.gov	▶ 🗑️ ▶▶	 Located	2019-01-07 12:56:52.708279-06:00	renlu23
432539	14716055@jobsub01.fnal.gov	▶ 🗑️ ▶▶			renlu23






Status History	
Start Timestamp	Status
2019-01-07 12:56:52	New
2019-01-07 13:04:33	Idle
2019-01-07 13:10:54	Running
2019-01-07 13:23:48	Held
2019-01-07 14:06:06	Running
2019-01-07 14:12:16	Held
2019-01-07 14:15:24	Completed
2019-01-07 14:21:04	Located

Stage submissions - Controlling, and monitoring real time state changes and viewing historical transitions are available.

POMS Landscape

POMS is well integrated with Landscape/Grafana with many plots, reports and status pages available.


Point in Time Shifters Report for [DUNE](#)



Active POMS Campaigns


Campaign ID	Campaign Name (click for stages)	Batch Jobs Completed	Failed Jobs	Average CPU Eff	Average Mem Eff
3572	hakan_RITM0848148_0	6	1	20%	24%
3463	np02_dc3_keepup_reco_skip300	2	0	98%	80%
3461	np02_dc3_keepup_reco_skip200	6	0	98%	89%
3460	np02_dc3_keepup_reco_skip150	3	0	99%	82%
3459	np02_dc3_keepup_reco_skip100	3	0	66%	91%
3458	np02_dc3_keepup_reco_skip50	9	1	99%	85%

POMS Clients



POMS a Successful Product

The success of POMS can be seen by its acceptance and continued usage by experiments.



- ★ Poms has managed over 9.8 million production jobs in the last year alone.
- ★ Usage has doubled in the last year.

Conclusion

- POMS simplifies organizing, tracking, and debugging your large-scale computing
- We continue to welcome feedback from DUNE towards new features, etc.