

Automated workflow logging

In order to improve debugging capabilities and declutter the screen when running auto_submit on a cybershake run, we wish to implement logging.

Logging levels

The default available levels of logging are as follows:

- Critical - messages relating to fatal issues
- Error - messages relating to serious, but not fatal issues
- Warning - possible issues
- Info - status reports for user information
- Debug - every significant action will have a debug message, these will not be printed to stdout by default though

Message format

The default logging message will be of the format

"%(asctime)s - %(module)s.%(funcName)s - %(message)s"

The logging message for tasks related to a realisation/process type pairing use the following format:

"%(levelname)8s -- %(asctime)s - %(module)s.%(funcName)s - {realisation}.{process_type} - %(message)s"

Where realisation and process_type are finalised at logger instantiation.

The following substitutions made by the python logger:

asctime	The time the message is logged. In the format YYYY-MM-DD HH:MM:SS,sss
module	The module the message is logged from
filename	The file the message is logged from
funcName	The function the message is logged from
message	The message being logged

Log file location

By default log files for auto_submit.py and queue_monitor.py will be placed in the cybershake root directory, with the name [calling script]_log_[timestamp].txt

This can be overridden by using the argument --log_file [file_location] when starting auto_submit or queue_monitor. This may be given as an absolute path, or relative to the cybershake root directory

Any existing files will be appended to and not deleted.

The end to end tests place these files in the test directory. --no_clean_up should be passed to it to preserve these files.

Usage

To add logging to a script in slurm_gm_workflow, shared_workflow.workflow_logger should be imported.

The function get_logger(name) should have a unique string passed to it. It will return a logger that will print messages with level greater than debug to stdout. It will also ignore any messages with a level beginning with NOPRINT, as defined in workflow_logger.

If file logging is desired then the logger should be passed to add_general_file_handler(logger, file_path).

get_basic_logger() will return a logger that prints messages with level greater than debug to stdout. It is intended to be used as a default if a function is called without passing a logger.

Examples of message sources

More concrete examples of what will be logged are as follows:

Level	Types of message
Critical	Invalid n_runs argument, duplicate db update file name, Unable to load workflow config
Error	Management db unknown status code found, HF/BB version invalid

Warning	Metadata storage passed invalid proc_type
Info	estimation model loading, tasks that are being skipped
Debug	Task submission, Dependency checking

- Dependency checking, task checking, job numbers etc
- Allow levels of log messages (i.e. DEBUG, INFO, WARNING, ERROR)
- Include timestamp and source (i.e. which module/function/class)

Example log output

Example output from a log file is available below. It was taken from the main auto submit thread of the run_cybershake automation wrapper.

```

DEBUG -- 2019-06-17 17:04:16,855 - MgmtDB._check_dependency_met - Considering task ProcessType.clean_up for
realisation AlpineF2K_HYP01-01_S1244. Completed tasks as follows: []
DEBUG -- 2019-06-17 17:04:16,855 - MgmtDB._check_dependency_met - (11, 'AlpineF2K_HYP01-01_S1244') has
remaining deps: [6]
INFO -- 2019-06-17 17:04:16,856 - auto_submit.run_main_submit_loop - Number of runnable tasks: 2006
INFO -- 2019-06-17 17:04:16,857 - auto_submit.run_main_submit_loop - Tasks to run this iteration:
Hossack_HYP240-501_S3634-EMOD3D, Hossack_HYP240-501_S3634-HF, Hossack_HYP235-501_S3584-EMOD3D, Hossack_HYP235-
501_S3584-HF, Hossack_HYP54-501_S1774-EMOD3D, Hossack_HYP54-501_S1774-HF, Hossack_HYP106-501_S2294-EMOD3D,
Hossack_HYP106-501_S2294-HF, Hossack_HYP275-501_S3984-EMOD3D, Hossack_HYP275-501_S3984-HF, Hossack_HYP118-
501_S2414-EMOD3D, Hossack_HYP118-501_S2414-HF
DEBUG -- 2019-06-17 17:04:16,858 - main auto submit - auto_submit.submit_task - Hossack_HYP240-501_S3634.
EMOD3D - Submit EMOD3D arguments: Namespace(account='nesi00213', auto=True, machine='maui', ncore=160, rel_dir='
/scale_wlg_nobackup/filesets/nobackup/nesi00213/RunFolder/jpa198/test/cybershake_test/version/Runs/Hossack
/Hossack_HYP240-501_S3634', retries=0, srf='Hossack_HYP240-501_S3634', write_directory='/scale_wlg_nobackup
/filesets/nobackup/nesi00213/RunFolder/jpa198/test/cybershake_test/version/Runs/Hossack/Hossack_HYP240-
501_S3634')
DEBUG -- 2019-06-17 17:04:16,873 - main auto submit - submit_emod3d.main - Hossack_HYP240-501_S3634.EMOD3D -
params.srf_file /nesi/nobackup/nesi00213/RunFolder/jpa198/test/cybershake_test/version/Data/Sources/Hossack/Srf
/Hossack_HYP240-501_S3634.srf
DEBUG -- 2019-06-17 17:04:16,873 - main auto submit - submit_emod3d.main - Hossack_HYP240-501_S3634.EMOD3D -
not set_params_only
DEBUG -- 2019-06-17 17:04:16,875 - main auto submit - submit_emod3d.main - Hossack_HYP240-501_S3634.EMOD3D -
retries has been set, but no check-pointing files exist. not scaling wct
DEBUG -- 2019-06-17 17:04:16,892 - main auto submit - set_runparams.create_run_params - Hossack_HYP240-
501_S3634.EMOD3D - emod3d_version not found as a key in e3d file. Ignoring variable. Value is 3.0.4.
DEBUG -- 2019-06-17 17:04:16,982 - main auto submit - auto_submit.submit_task - Hossack_HYP240-501_S3634.HF
- Submit HF arguments: Namespace(account='nesi00213', auto=True, debug=False, machine='maui', ncore=80,
rel_dir='/scale_wlg_nobackup/filesets/nobackup/nesi00213/RunFolder/jpa198/test/cybershake_test/version/Runs
/Hossack/Hossack_HYP240-501_S3634', retries=0, seed=0, site_specific=None, srf='Hossack_HYP240-501_S3634',
version='run_hf_mpi', write_directory='/scale_wlg_nobackup/filesets/nobackup/nesi00213/RunFolder/jpa198/test
/cybershake_test/version/Runs/Hossack/Hossack_HYP240-501_S3634')
DEBUG -- 2019-06-17 17:04:16,992 - main auto submit - submit_hf.main - Hossack_HYP240-501_S3634.HF -
version: run_hf_mpi
DEBUG -- 2019-06-17 17:04:16,994 - main auto submit - model.predict - Hossack_HYP240-501_S3634.HF - The
entry is out of bounds. The SVR models will be used for estimation.
DEBUG -- 2019-06-17 17:04:17,010 - main auto submit - shared.submit_sl_script - Hossack_HYP240-501_S3634.HF
- Submitting /scale_wlg_nobackup/filesets/nobackup/nesi00213/RunFolder/jpa198/test/cybershake_test/version/Runs
/Hossack/Hossack_HYP240-501_S3634/run_hf_mpi_Hossack_HYP240-501_S3634_20190617_170417.sl on machine maui
DEBUG -- 2019-06-17 17:04:17,077 - main auto submit - auto_submit.submit_task - Hossack_HYP235-501_S3584.
EMOD3D - Submit EMOD3D arguments: Namespace(account='nesi00213', auto=True, machine='maui', ncore=160, rel_dir='
/scale_wlg_nobackup/filesets/nobackup/nesi00213/RunFolder/jpa198/test/cybershake_test/version/Runs/Hossack
/Hossack_HYP235-501_S3584', retries=0, srf='Hossack_HYP235-501_S3584', write_directory='/scale_wlg_nobackup
/filesets/nobackup/nesi00213/RunFolder/jpa198/test/cybershake_test/version/Runs/Hossack/Hossack_HYP235-
501_S3584')
DEBUG -- 2019-06-17 17:04:17,088 - main auto submit - submit_emod3d.main - Hossack_HYP235-501_S3584.EMOD3D -
params.srf_file /nesi/nobackup/nesi00213/RunFolder/jpa198/test/cybershake_test/version/Data/Sources/Hossack/Srf
/Hossack_HYP235-501_S3584.srf
DEBUG -- 2019-06-17 17:04:17,088 - main auto submit - submit_emod3d.main - Hossack_HYP235-501_S3584.EMOD3D -
not set_params_only
DEBUG -- 2019-06-17 17:04:17,089 - main auto submit - submit_emod3d.main - Hossack_HYP235-501_S3584.EMOD3D -
retries has been set, but no check-pointing files exist. not scaling wct
DEBUG -- 2019-06-17 17:04:17,105 - main auto submit - set_runparams.create_run_params - Hossack_HYP235-
501_S3584.EMOD3D - emod3d_version not found as a key in e3d file. Ignoring variable. Value is 3.0.4.
DEBUG -- 2019-06-17 17:04:17,233 - main auto submit - auto_submit.submit_task - Hossack_HYP235-501_S3584.HF
- Submit HF arguments: Namespace(account='nesi00213', auto=True, debug=False, machine='maui', ncore=80,
```

```
rel_dir='/scale_wlg_nobackup/filesets/nobackup/nesi00213/RunFolder/jpa198/test/cybershake_test/version/Runs/Hossack/Hossack_HYP235-501_S3584', retries=0, seed=0, site_specific=None, srf='Hossack_HYP235-501_S3584', version='run_hf_mpi', write_directory='/scale_wlg_nobackup/filesets/nobackup/nesi00213/RunFolder/jpa198/test/cybershake_test/version/Runs/Hossack/Hossack_HYP235-501_S3584')
DEBUG -- 2019-06-17 17:04:17,243 - main auto submit - submit_hf.main - Hossack_HYP235-501_S3584.HF -
version: run_hf_mpi
DEBUG -- 2019-06-17 17:04:17,245 - main auto submit - model.predict - Hossack_HYP235-501_S3584.HF - The entry is out of bounds. The SVR models will be used for estimation.
DEBUG -- 2019-06-17 17:04:17,260 - main auto submit - shared.submit_s1_script - Hossack_HYP235-501_S3584.HF
- Submitting /scale_wlg_nobackup/filesets/nobackup/nesi00213/RunFolder/jpa198/test/cybershake_test/version/Runs/Hossack/Hossack_HYP235-501_S3584/run_hf_mpi_Hossack_HYP235-501_S3584_20190617_170417.s1 on machine maui
DEBUG -- 2019-06-17 17:04:17,323 - main auto submit - auto_submit.submit_task - Hossack_HYP54-501_S1774.
EMOD3D - Submit EMOD3D arguments: Namespace(account='nesi00213', auto=True, machine='maui', ncore=160, rel_dir='/scale_wlg_nobackup/filesets/nobackup/nesi00213/RunFolder/jpa198/test/cybershake_test/version/Runs/Hossack/Hossack_HYP54-501_S1774', retries=0, srf='Hossack_HYP54-501_S1774', write_directory='/scale_wlg_nobackup/filesets/nobackup/nesi00213/RunFolder/jpa198/test/cybershake_test/version/Runs/Hossack/Hossack_HYP54-501_S1774')
DEBUG -- 2019-06-17 17:04:17,334 - main auto submit - submit_emod3d.main - Hossack_HYP54-501_S1774.EMOD3D - params.srf_file /nesi/nobackup/nesi00213/RunFolder/jpa198/test/cybershake_test/version/Data/Sources/Hossack/Srf/Hossack_HYP54-501_S1774.srf
DEBUG -- 2019-06-17 17:04:17,334 - main auto submit - submit_emod3d.main - Hossack_HYP54-501_S1774.EMOD3D - not set_params_only
DEBUG -- 2019-06-17 17:04:17,335 - main auto submit - submit_emod3d.main - Hossack_HYP54-501_S1774.EMOD3D - retries has been set, but no check-pointing files exist. not scaling wct
```