Release Notes: National Water Model Version 3.0.0

List of Major Enhancements in Version 3.0.0:

- First time provision of NWM Total Water Level guidance for coastal areas of the Continental United States (CONUS), Hawaii and Puerto Rico / US Virgin Island domains. This is accomplished via use of the Semi-implicit Cross-scale Hydroscience Integrated System Model (SCHISM) integrated within the NWM, to couple NWM freshwater discharge estimates with oceanic forcing from the Surge and Tide Operational Forecast System (STOFS) and Probabilistic Tropical Storm Surge (P-SURGE) model.
- NWM Domain expansion to south-central Alaska (Cook Inlet, Copper River Basin, and Prince William Sound regions), enabling provision of NWM operational hydrologic model forecast guidance to this region.
- Addition of the National Blend of Models (NBM) as a forcing source for NWM CONUS medium-range forecasts and Alaska short-range and medium-range forecasts.
- Use of Multi-Radar Multi-Sensor (MRMS) precipitation as forcing for the NWM Analysis and Assimilation configuration over the Puerto Rico / US Virgin Island domain.
- Ingest of RFC-supplied reservoir outflow forecasts at 57 additional locations, bringing the total of such sites to 392.
- Enhancements to the treatment of reservoirs, land surface parameters and calibration/regionalization approach leading to improvements in model skill.
- Implementation of a new surface runoff scheme (Xinanjiang) within the NWM, improving streamflow simulation performance.
- Various hydrofabric improvements for all domains
- Removal of pre-conversion (grib2netcdf) configurations
- Removal of the NWM medium-range member 7 configuration
- Addition of the ‘restart’ capability for models that have runtimes of more than 1 hour, namely, the medium range members 1 to 6, long range members 1 to 4, coastal medium ranges, coastal medium range blend configurations, and coastal medium range with psurge configurations. This feature allows a model to restart at a point at which it has stopped or crashed without starting from the very beginning of the simulation. It will save time in case there is a failure and the re-run is needed in realtime for jobs with long runtimes.

- Compute resource information
    - Job type (serial or parallel) and associated node use
        - Serial
            - USGS, CANADA, MERGE_USGS_CA, USACE and RFC streamflow observation processing (1 node for each job using 4GB - 24GB of memory)
            - FORCING_MPE (1 node, 1 core using 20GB)

- Parallel
    - Forcing engine jobs for analysis and assimilation, short-range, medium-range and long-range configurations , including CONUS, Alaska, Hawaii and Puerto Rico domains (1 node for each job using 10GB-512GB of memory)
    - Model jobs for analysis and assimilation, short-range, medium-range and long-range configurations, including CONUS, Alaska, Hawaii and Puerto Rico domains (1 to 8 nodes for each job, using 5GB - 512GB memory)
    - Model open-loop jobs for analysis and assimilation (normal, long and extended) and medium-range configurations including CONUS, Alaska, Hawaii and Puerto Rico domains (1 to 8 nodes for each job, using 12GB - 100 GB memory)
    - Coastal Model jobs for analysis_assimilation, short-range, medium-range, including Atlgulf, Pacific, Hawaii and Puertorico domains (exclhost 1 to 8 nodes for each job, using 512 GB memory)

- Runtimes (on WCOSS)
    - USGS, CANADA, MERGE_USGS_CA, USACE and RFC streamflow observation processing
        - JNWM_USGS_TIMESLICES: ~ 6.5 minutes (1 node - 1 core per node)
        - JNWM_ACE_TIMESLICES: ~ 1 minutes (1 node - 1 core)
        - JNWM_CANADA_TIMESLICES: ~ 1.5 minutes (1 node - 1 core)
        - JNWM_MERGE_USGS_CA_TIMESLICES: ~ 1 minutes (1 node - 1 core)
        - JNWM_RFC_TIMESLICES: ~ 6.5 minutes (1 node - 1 core)

    - Forcing engine
        - JNWM_FORCING_ANALYSIS_ASSIM: ~ *3 minutes* (1 node 72 cores per node)
        - JNWM_FORCING_MPE: ~ *3.5 minutes* (1 node, 1 core)
        - JNWM_FORCING_ANALYSIS_ASSIM_EXTEND: ~ *1.6 minutes* (1 nodes - 14 cores per node)
        - JNWM_FORCING_SHORT_RANGE: ~ *3 minutes* (exclhost 1 node - 128 cores per node)
        - JNWM_FORCING_MEDIUM_RANGE: ~ *13.1 minutes* (exclhost 1 node - 128 cores per node)
        - JNWM_FORCING_MEDIUM_RANGE_BLEND: ~ *17.2 minutes* (exclhost 1 node - 128 cores per node)
        - JNWM_FORCING_LONG_RANGE (Long Range 1-4): ~ *14.1 minutes* (exclhost 1 node - 128 cores per node)

- JNWM_FORCING_ANALYSIS_ASSIM_HAWAII: ~ *1 minutes* (1 node - 8 cores per node)
- JNWM_FORCING_SHORT_RANGE_HAWAII: ~ *1 minute* (1 node - 4 cores per node)
- JNWM_FORCING_ANALYSIS_ASSIM_PUERTORICO: ~ *1 minute* (1 node - 8 cores per node)
- JNWM_FORCING_SHORT_RANGE_PUERTORICO: ~ *1 minute* (1 node - 4 cores per node)
- JNWM_FORCING_ANALYSIS_ASSIM_ALASKA: ~ 1 *minutes* (1 node - 8 cores per node)
- JNWM_FORCING_ANALYSIS_ASSIM_EXTEND_ALASKA: ~ 1 *minutes* (1 nodes - 16 cores per node)
- JNWM_FORCING_SHORT_RANGE_ALASKA: ~ 1.1 *minute* (1 node - 4 cores per node)
- JNWM_FORCING_MEDIUM_RANGE_ALASKA: ~ 5.35 *minutes* (exclhost 1 node - 128 cores per node)
- JNWM_FORCING_MEDIUM_RANGE_BLEND_ALASKA: ~ 7.4 *minutes* (exclhost 1 node - 128 cores per node)

- Model runs
  - JNWM_MODEL (Analysis and Assimilation): ~ 10 *minutes* (exclhost 6 nodes - 128 cores per node)
  - JNWM_MODEL (Extended Analysis and assimilation): ~ *27 minutes* (exclhost 6 nodes - 128 cores per node)
  - JNWM_MODEL (Short Range): ~ 16 *minutes* (exclhost 6 nodes - 128 cores per node)
  - JNWM_MODEL (Medium Range 1-6): ~ *87 minutes* (exclhost 8 nodes - 128 cores per node)
  - JNWM_MODEL (Medium Range Blend): ~ 88 *minutes* (exclhost 8 nodes - 128 cores per node)
  - JNWM_MODEL (Long Range 1-4): ~ *54 minutes* (exclhost 6 nodes - 128 cores per node)
  - JNWM_MODEL (Long-range Analysis and assimilation): ~ *8.3 minutes* (exclhost 1 nodes - 128 cores per node)
  - JNWM_MODEL (Hawaii Analysis and assimilation): ~ 4.1 *minutes* (1 node - 9 cores per node)
  - JNWM_MODEL (Hawaii Short Range): ~13.7 *minutes* (exclhost 1 node - 128 cores per node)
  - JNWM_MODEL (PuertoRico Analysis and assimilation): ~ *1 minutes* (1 node - 9 cores per node)
  - JNWM_MODEL (PuertoRico Short Range): ~ 7.28 *minutes* (1 node - 9 cores per node)

- JNWM_MODEL (Alaska Analysis and assimilation): ~ 2.1 *minutes* (1 node - 9 cores per node)
- JNWM_MODEL (Alaska Extended Analysis and assimilation): ~ 3.1 *minutes* (exclhost 1 node - 128 cores per node)
- JNWM_MODEL (Alaska Short Range): ~ *7.6 minutes* (1 node - 24 cores per node)
- JNWM_MODEL (Alaska Medium Range 1-6): ~ 11 *minutes* (exclhost 1 node - 128 cores per node)
- JNWM_MODEL (Alaska Medium Range Blend): ~ 13 *minutes* (exclhost 1 node - 128 cores per node)

■ Open-loop (non-DA) model runs
- NWM_MODEL_NO_DA (Analysis and Assimilation): ~ *1 minutes* (6 nodes - 24 cores per node)
- NWM_MODEL_NO_DA (Extended Analysis and Assimilation): ~ *3 minutes* (6 nodes - 28 cores per node)
- NWM_MODEL_NO_DA (Hawaii Analysis and Assimilation): ~ 1 *minute* (1 node - 9 cores per node)
- NWM_MODEL_NO_DA (PuertoRico Analysis and Assimilation): ~ *1 minute* (1 node - 9 cores per node)
- NWM_MODEL_NO_DA (Long-range Analysis and Assimilation): ~ *3 minutes* (exclhost 1 node - 128 cores per node)
- NWM_MODEL_NO_DA (Hawaii Short-range): ~ *2 minutes* (1 node - 9 cores per node)
- NWM_MODEL_NO_DA (PuertoRico Short-range): ~ *1 minute* (1 node - 9 cores per node)
- NWM_MODEL_NO_DA (Medium-range): ~ 12 *minutes* (8 nodes - 64 cores per node)
- NWM_MODEL_NO_DA (Alaska Analysis and Assimilation): ~ *minutes* (1 node - 9 cores per node)
- NWM_MODEL_NO_DA (Alaska Extended Analysis and Assimilation): ~ 1 *minutes* (1 node - 128 cores per node)
- NWM_MODEL_NO_DA (Alaska Medium-range): ~ 9 *minutes* (1 node - 24 cores per node)

■ Coastal model runs
- JNWM_COASTAL (Atlgulf Analysis and Assimilation): ~ 18 *minutes* (exclhost 6 nodes - 128 cores per node)
- JNWM_COASTAL (Atlgulf Extended Analysis and assimilation): ~ *30 minutes* (exclhost 6 nodes - 128 cores per node)
- JNWM_COASTAL (Atlgulf Short Range): ~ 25 *minutes* (exclhost 6 nodes - 128 cores per node)

- JNWM_COASTAL (Atlgulf Medium Range1): ~ *6 hrs 34 minutes* (exclhost 8 nodes - 128 cores per node)
- JNWM_COASTAL (Atlgulf Medium Range Blend): ~ *6 hrs 24 minutes* (8 nodes - 128 cores per node)
- JNWM_COASTAL (***Psurge*** Atlgulf Short Range): ~ *25 minutes* (exclhost 6 nodes - 128 cores per node)
- JNWM_COASTAL (***Psurge*** Atlgulf Medium Range1): ~ *5 hrs 50 minutes* (exclhost 8 nodes - 128 cores per node)
- JNWM_COASTAL (***Psurge*** Atlgulf Medium Range Blend): ~ *5 hrs 38 minutes* (exclhost 8 nodes - 128 cores per node)
- JNWM_COASTAL (Pacific Analysis and assimilation): ~ 8 *minutes* (exclhost 2 nodes - 128 cores per node)
- JNWM_COASTAL (Pacific Extended Analysis and assimilation): ~ 22 *minutes* (exclhost 2 nodes - 128 cores per node)
- JNWM_COASTAL (Pacific Short Range): ~ 13 *minutes* (exclhost 4 nodes - 128 cores per node)
- JNWM_COASTAL (Pacific Medium Range 1-6): ~ 3 hrs 26 *minutes* (exclhost 8 nodes - 128 cores per node)
- JNWM_COASTAL (Pacific Medium Range Blend): ~ 3 hrs 18 *minutes* (exclhost 8 nodes - 128 cores per node)
- JNWM_COASTAL (Hawaii Analysis and assimilation): ~ *3 minutes* (exclhost 1 node - 128 cores per node)
- JNWM_COASTAL (Hawaii Short Range): ~ *13 minutes* (exclhost 4 node - 128 cores per node)
- JNWM_COASTAL (PuertoRico Analysis and assimilation): ~ *9 minutes* (exclhost 1 node - 128 cores per node)
- JNWM_COASTAL (PuertoRico Short Range): ~ 37 *minutes* (exclhost 2 nodes - 128 cores per node)

- Disk space required per day
  - COM area: ~ 6.3 Tb/day of forcing, model output files and intermediate data (nwges) (but only ~1.2 TB is transferred to NOMADS each day)
  - COM/nwges area for recent files: ~3.4 Tb/day (3.4Tb nwm; 3.2Gb nwm_forcing)

- Cycling frequency
  - Every 15 minutes: USGS, ACE, CANADA, RFC streamflow observation processing
  - Hourly: Conus, Hawaii, PuertoRico and Alaska analysis and assimilation forcing engine,model and coastal; Conus short range forcing engine, model and coastal.
  - Every 3 hours: 00z,03z..21z Alaska Short-Range forcing and model.

        - Every 6 hours: long range analysis and assimilation, long-range (4 members execute four times per day for a total of 16 ensemble members valid for the same time period) forcing engine and model; medium-range forcing, model (6 members execute at the same time period) and coastal; blend medium-range forcing, model and coastal.
        - Every 12 hours: 00z and 12z Hawaii Short-range forcing engine and model; 06z and 18z PuertoRico Short-Range forcing engine and model.
        - Daily: MPE forcing; Extended analysis assimilation forcing, model and coastal (executed 1 time per day)

    - Version of libs/compilers/shared code being used
        - Intel IFORT 19.1.3.304
        - Intel gcc 11.2.0
        - IOBUF 2.0.10
        - netcdf/4.7.4
        - Bash 4.4.23
        - hdf5/1.10.6
        - wgrib2/2.0.8
        - python/3.8.6 (python packages: mpi4py 3.1.0, ESMPy 8.0.0, netCDF4 1.5.6, numpy 1.20.1)

    - Data retention
        - Files in /com should be retained for 10 days
        - Files on NOMADS should be retained for 48 hours per file (not per directory)

- Pre-implementation testing requirements
    - The following jobs should be tested
        - Streamflow observation processing (JNWM_USGS_TIMESLICES, JNWM_ACE_TIMESLICES, JNWM_CANADA_TIMESLICES, JNWM_MERGE_USGS_CA_TIMESLICES, and JNWM_RFC_TIMESERIES)
        - Conus, Hawaii, PuertoRico and Alaska Analysis and assimilation forcing engine (JNWM_FORCING_ANALYSIS_ASSIM, NWM_FORCING_ANALYSIS_ASSIM_HAWAII, NWM_FORCING_ANALYSIS_ASSIM_PUERTORICO and NWM_FORCING_ANALYSIS_ASSIM_ALASKA)
        - Conus, Hawaii, PuertoRico and Alaska Short-range forcing engine (JNWM_FORCING_SHORT_RANGE, JNWM_FORCING_SHORT_RANGE_HAWAII, JNWM_FORCING_SHORT_RANGE_PUERTORICO, and JNWM_FORCING_SHORT_RANGE_ALASKA)

- Conus and Alaska Medium-range forcing engine (JNWM_FORCING_MEDIUM_RANGE, JNWM_FORCING_MEDIUM_RANGE_BLEND, JNWM_FORCING_MEDIUM_RANGE_ALASKA, JNWM_FORCING_MEDIUM_RANGE_BLEND_ALASKA)
- Long-range forcing engine (JNWM_FORCING_LONG_RANGE)
- Extended forcing (JNWM_FORCING_MPE,
- JNWM_FORCING_ANALYSIS_ASSIM_EXTEND)

- Conus, Hawaii, PuertoRico, and Alaska Analysis and Assimilation model run (JNWM_MODEL)
- Conus and Alaska Extended Analysis and Assimilation model run (JNWM_MODEL)
- Long-range Analysis and Assimilation model run (JNWM_MODEL)
- Conus, Hawaii, PuertoRico, and Alaska Short-range model run (JNWM_MODEL)
- Conus and Alaska Medium-range and Medium-range Blend model run (JNWM_MODEL)
- Long-range model run (JNWM_MODEL)

- Open-loop with non-DA model runs: Conus, Hawaii, PuertoRico and Alaska Analysis and Assimilation; Long Analysis and Assimilation, Conus and Alaska Extended AnA; Hawaii and PuertoRico Short-range and Conus and Alaska Medium_Range (JNWM_MODEL_NO_DA)

- Coastal model run:
    - Atlgulf, Pacific, Hawaii, and PuertoRico Analysis and Assimilation model run (JNWM_COASTAL)
    - Atlgulf and Pacific Extended Analysis and Assimilation model run (JNWM_COASTAL)
    - Atlgulf, Pacific, Hawaii, and PuertoRico Short-range model runs (JNWM_COASTAL)
    - Atlgulf and Pacific Medium-range and Medium-range Blend model run (JNWM_COASTAL)
    - **Psurge Atlgulf** Short-range, Medium-range, and Medium-range Blend model run (JNWM_COASTAL)

List of Enhancements in Version 2.2.0:

Ported from WCOSS1 to WCOSS2

- Updated modules, compiler and libraries to build the model binary.
- Updated version numbers for modules, compiler and libraries.
- Replaced LFS directives with PBS directives in the EcFlow scripts.

- Added run.ver and build.ver files in the versions directory.
- Updated build and installation scripts to be compatible with the WCOSS2 environment.
- Updated the ecFlow, jobs, and ex-scripts as required by the NCO transition checklist, such as $COMOUT, $DCOM, $NWGES environmental variables and removed jlogfile, etc.
- Added testing scripts to facilitate testing using the canned data NCO provided.
- Ran timing tests for all configurations to optimize resource usage. Roughly optimized number of nodes and cores along with memory allocation. Achieved faster runtimes on WCOSS2 than were seen with NWM v2.1 on WCOSS1, while using similar or lower compute resources.
- Reduced the extremely long lines in the CFP command files for the long range forcing configuration. The CFP crashes when the command line is too long.
- Note that all parameter files remain unchanged from v2.1 (identical parameter files are used for v2.2 as changes were not necessary for porting functionality)

List of Enhancements in Version 2.1.0:

- Addition of new Puerto Rico / US Virgin Islands domain
- Addition of new water management (reservoir) module
- Expansion of CONUS streamflow routing domain to include Great Lakes drainage region
- Inclusion of 8 Field-requested open loop (non-DA) job configurations:
  - 5 no-DA Analysis configs (no-DA versions of existing AnA configs, AnA, AnA Extend, AnA Long Range, Hawaii AnA, and Puerto Rico AnA)
  - 3 no-DA forecast configs (CONUS MR, Hawaii SR, Puerto Rico SR)
- NWM.v2.0 triggers - Google Slides Modification of Hawaii Short Range forecast (reduced to twice per day out to 48 hours but augmented with improved forcing per Field request)
- Ingest of new data sources
  - WRF-ARW for Puerto Rico and Hawaii (now used with NAM-Nest)
  - Canadian streamflow observations, USACE stream/reservoir observations, and RFC reservoir forecasts
  - MRMS v12 QPE for CONUS and Hawaii
- Updated Forcing Engine, converted from NCL code-base to Python-base
- Addition of pre-conversion jobs to convert forcing inputs on-the-fly (efficiency boost)
- Upgraded snow pack and surface water physics
- Enhanced calibration and upgrade of land surface vegetation parameter data set

- Compute resource information
  - Job type (serial or parallel) and associated node use
    - Serial

        - USGS, CANADA, MERGE_USGS_CA, USACE and RFC streamflow observation processing (1 node)
        - Pre-conversion forcing data (1 node)

    - Parallel
        - Forcing engine jobs for analysis and assimilation, short-range, medium-range and long-range configurations (1-20 nodes for each job using 500MB-3000MB of memory)
        - Model jobs for analysis and assimilation, short-range, medium-range and long-range configurations, including CONUS, Hawaii and Puerto Rico domains (1 to 32 nodes for each job, using 100 MB memory)
        - Model open-loop jobs for analysis and assimilation (normal, long and extended) and medium-range configurations including CONUS, Hawaii and Puerto Rico domains (1 to 32 nodes for each job, using 100 MB memory)

- Runtimes(on WCOSS)
    - USGS, CANADA, MERGE_USGS_CA, USACE and RFC streamflow observation processing
        - JNWM_USGS_TIMESLICES: ~ 4 minutes (1 node - 24 cores per node)
        - JNWM_ACE_TIMESLICES: ~ 1 minutes (1 node - 24 cores)
        - JNWM_CANADA_TIMESLICES: ~ 23 minutes (1 node - 24 cores)
        - JNWM_MERGE_USGS_CA_TIMESLICES: ~ 2.5 minutes (1 node - 24 cores)
        - JNWM_RFC_TIMESLICES: ~ 11 minutes (1 node - 24 cores)

    - Pre-conversion forcing data (time below are for a full cycle...will be much shorter if timed on a per-file basis)
        - JNWM_GRIB2NETCDF(cfs): ~ *11.9 minutes* (1 nodes - 1 core)

- JNWM_GRIB2NETCDF(gfs): *~ 12.3 minutes* (1 nodes - 1 core)
- JNWM_GRIB2NETCDF(hiresw): *~ 3 minutes* (1 nodes - 1 core)
- JNWM_GRIB2NETCDF(hrrr): *~ 3 minutes* (1 nodes - 1 core)
- JNWM_GRIB2NETCDF(multisensor_qpe): *~ 2 minute* (1 nodes - 1 core)
- JNWM_GRIB2NETCDF(nam): *~ 6 minutes* (1 nodes - 1 core)
- JNWM_GRIB2NETCDF(radaronly): *~ 2 minutes* (1 nodes - 1 core)
- JNWM_GRIB2NETCDF(rap): *~ 4 minutes* (1 nodes - 1 core)

■ Forcing engine
- JNWM_FORCING_ANALYSIS_ASSIM: *~ 3 minutes* (4 node 18 cores per node)
- JNWM_FORCING_SHORT_RANGE: *~ 4 minutes* (4 nodes - 18 cores per node)
- JNWM_FORCING_MEDIUM_RANGE: *~ 21 minutes* (20 nodes - 12 cores per node)
- JNWM_FORCING_LONG_RANGE (Long Range 1-4): *~ 28 minutes* (2 nodes - 24 cores per node)
- JNWM_FORCING_MPE: *~ 12 minutes* (1 node - 1 core)
- JNWM_FORCING_ANALYSIS_ASSIM_EXTEND: *~ 3 minutes* (2 nodes - 14 cores per node)
- JNWM_FORCING_ANALYSIS_ASSIM_HAWAII: *~ 2 minutes* (1 node - 24 cores per node)
- JNWM_FORCING_SHORT_RANGE_HAWAII: *~ 1 minute* (4 nodes - 16 cores per node)
- JNWM_FORCING_ANALYSIS_ASSIM_PUERTORICO: *~ 1 minute* (1 node - 24 cores per node)
- JNWM_FORCING_SHORT_RANGE_PUERTORICO: *~ 1 minute* (5 nodes - 10 cores per node)

- Model runs
    - JNWM_MODEL (Analysis and Assimilation): ~ *10 minutes* (32 nodes - 24 cores per node)
    - JNWM_MODEL (Short Range): ~ *14 minutes* (32 nodes - 24 cores per node)
    - JNWM_MODEL (Medium Range 1-7): ~ *80 minutes* (32 nodes - 24 cores per node)
    - JNWM_MODEL (Long Range 1-4): ~ *92 minutes* (32 nodes - 24 cores per node)
    - JNWM_MODEL (Extended Analysis and assimilation): ~ *27 minutes* (32 nodes - 24 cores per node)
    - JNWM_MODEL (Long-range Analysis and assimilation): ~ *12 minutes* (20 nodes - 24 cores per node)
    - JNWM_MODEL (Hawaii Analysis and assimilation): ~ *3.7 minutes* (1 node - 9 cores per node)
    - JNWM_MODEL (Hawaii Short Range): ~ *31.7 minutes* (1 node - 9 cores per node)
    - JNWM_MODEL (PuertoRico Analysis and assimilation): ~ *2 minutes* (1 node - 9 cores per node)
    - JNWM_MODEL (PuertoRico Short Range): ~ *7 minutes* (1 node - 9 cores per node)

- Open-loop (non-DA) model runs
    - NWM_MODEL_NO_DA (Analysis and Assimilation): ~ *2 minutes* (18 nodes - 24 cores per node)
    - NWM_MODEL_NO_DA (Extended Analysis and Assimilation): ~ *6 minutes* (32 nodes - 24 cores per node)
    - NWM_MODEL_NO_DA (Hawaii Analysis and Assimilation): ~ *1 minute* (1 node - 9 cores per node)
    - NWM_MODEL_NO_DA (PuertoRico Analysis and Assimilation): ~ *1 minute* (1 node - 9 cores per node)
    - NWM_MODEL_NO_DA (Long-range Analysis and Assimilation): ~ *3 minutes* (20 nodes - 24 cores per node)

- NWM_MODEL_NO_DA (Hawaii Short-range): ~ *2 minutes* (1 node - 9 cores per node)
- NWM_MODEL_NO_DA (PuertoRico Short-range): ~ *1 minute* (1 node - 9 cores per node)
- NWM_MODEL_NO_DA (Medium-range): ~ *38.6 minutes* (4 nodes - 24 cores per node)

- Disk space required per day
  - COM area: ~ 2.4 Tb/day of forcing and model output files (but only ~915 GB is transferred to NOMADS each day)
  - NWGES area for recent files: ~3.1 Tb/day (2.9Tb nwm; 3.2Gb nwm_forcing; and 196Gb nwm.grib2netcdf)
- Cycling frequency
  - Every 15 minutes: USGS, ACE, CANADA, RFC streamflow observation processing
  - Hourly: Conus, Hawaii and PuertoRico analysis and assimilation forcing engine and model; Conus short range forcing engine and model.
  - Every 6 hours: long range analysis and assimilation, long-range (4 members execute four times per day for a total of 16 ensemble members valid for same time period) forcing engine and model; medium-range forcing and model (7 members execute at the same time period).
  - Every 12 hours: 00z and 12z Hawaii Short-range forcing engine and model; 06z and 18z PuertoRico Short-Range forcing engine and model.
  - Daily: MPE forcing; Extended analysis assimilation forcing and model (executed 1 time per day)

- Version of libs/compilers/shared code being used
  - Intel Fortran 18.1.163
  - IOBUF 2.0.5
  - Hugepages
  - cray-netcdf 4.3.2 and NetCDF 4.7.4
  - Python 3.6.3
  - Python NetCDF4 package
  - Bash 4.2.53
  - HDF5 1.8.13 and 1.10.6

        - mpi4py/3.0.2
        - ESMF-gnu-sandybridge
        - ESMPy/7_1_0
        - grib_util/1.0.5

    - Data retention
        - Files in /com should be retained for 10 days
        - Files in /nwges should be retained for 5 days
        - Files on NOMADS should be retained for 48 hours per file (not per directory) **Note that this is an update from V2.0**

- Pre-implementation testing requirements
    - The following jobs should be tested
        - Streamflow observation processing (JNWM_USGS_TIMESLICES, JNWM_ACE_TIMESLICES, JNWM_CANADA_TIMESLICES, JNWM_MERGE_USGS_CA_TIMESLICES, and JNWM_RFC_TIMESERIES)

        - Pre-conversion forcing GFS, CFS, RAP, HRRR, HIRESW, NAM-nest, RADARONLY, Conus and Hawaii MUILTISENSOR QPE (JNWM_GRIB2NETCDF)

        - Conus, Hawaii and PuertoRico Analysis and assimilation forcing engine (JNWM_FORCING_ANALYSIS_ASSIM, NWM_FORCING_ANALYSIS_ASSIM_HAWAII, and NWM_FORCING_ANALYSIS_ASSIM_PUERTORICO)
        - Conus, Hawaii, and PuertoRico Short-range forcing engine (JNWM_FORCING_SHORT_RANGE, JNWM_FORCING_SHORT_RANGE_HAWAII, and JNWM_FORCING_SHORT_RANGE_PUERTORICO)
        - Medium-range forcing engine (JNWM_FORCING_MEDIUM_RANGE)
        - Long-range forcing engine (JNWM_FORCING_LONG_RANGE)
        - Extended forcing (JNWM_FORCING_MPE,

JNWM_FORCING_ANALYSIS_ASSIM_EXTEND)

- Conus, Hawaii, and PuertoRico Analysis and Assimilation model run (JNWM_MODEL)
- Extended Analysis and Assimilation model run (JNWM_MODEL)
- Long-range Analysis and Assimilation model run (JNWM_MODEL)
- Conus, Hawaii, and PuertoRico Short-range model run (JNWM_MODEL)
- Medium-range model run (JNWM_MODEL)
- Long-range model run (JNWM_MODEL)

- Open-loop with non-DA model runs: Conus, Hawaii and PuertoRico Analysis and Assimilation; Long Analysis and Assimilation, Extended AnA; Hawaii and PuertoRico Short-range and Medium_Range (JNWM_MODEL_NO_DA)

List of Enhancements in Version 2.0.1:
- Modified MPE forcing code to read in the new upstream Stage IV filename and grib2 data format.

List of Enhancements in Version 2.0.0:
-Addition of an Extended Analysis configuration (daily 28-hour look-back using RFC-based MPE precipitation from NCEP Stage IV dataset)
-Addition of Hawaii to NWM domain (including 3-hr Analysis and 60-hr Short-Range forecast--both forced by the NAM-Nest NWP model)
-Addition of a separate Long-Range Analysis configuration to initialize the Long-Range forecast
-Addition of a Medium Range ensemble forecast configuration (7 members 4 x day) (mem1=uses current GFS to 10 days, mem2-7=use time lagged GFS out to 8.5 days)
-Use of 13km GFS forcing (versus 0.25 degree in NWM V1.2)
-Improved downscaling of GFS and CFS forcing via RFC Mountain Mapper-based approach
-Improved physics (out-of-bank parameterization via compound channel, improved snow physics)
-Improved and expanded calibration of hydrologic parameters
-Corrections to stream connectivity
-Improved code modularity
-Enhanced warning notification system by adding an email capability--Important warning messages will be sent as appropriate to the SDM, SPAs, OWP developers or all of them indicating any job issues which occur.
-Refined land surface and hydrologic parameters by expanding calibration from ~1100 to ~1400 calibration basins and improving parameter regionalization process.

- Various hydrofabric improvements including:
    - Fixes to 37 stream breaks
    - Addition of 13,637 new flowlines (Hawaii stream reaches)
    - Addition of OCONUS basins in the Hawaii domain (16,625 km^2)
    - Addition of 58 USGS stream gauges into assimilation routine (Hawaii domain)
    - Addition of 3,955 CONUS reservoirs (now totaling 5,461)
    - Addition of 10 new reservoirs in the Hawaii domain.
    - Inclusion of a new elevation base that is harmonized with the NHDPlus channel network

-Changes to the fields and metadata contained within NWM output:

The following new variables are added to the existing Forcing, Analysis and Assimilation, Short-, Medium- and Long-Range forecast configurations, and are also present in the new Forcing, Hawaii Short-Range, Hawaii Analysis, Extended Analysis and Long-Range Analysis configurations:

- GDAL-style “crs” variable that provides information on the Coordinate Reference Systems
- "elevation" is now "Feature Elevation"
- "order" is now "Streamflow Order"
- “time” with attributes of “valid_min” and “valid_max”

The following new variables are added to the preexisting Analysis, and new Extended and Long-Range Analysis and Assimilation configurations:

- “SNLIQ” is snow layer liquid water (units: mm)
- “ISNOW” is number of snow layers (units: count)
- “SOIL_M” is volumetric soil moisture (units: m3/m3)
- “SOILICE” is fraction of soil moisture that is ice (units: 1)
- “SOIL_T” soil temperature (units: K)
- Compute resource information
    - Job type (serial or parallel) and associated node use
        - Serial
            - USGS streamflow observation processing (1 node)
        - Parallel
            - Forcing engine jobs for analysis and assimilation, short-range, medium-range and long-range configurations (1-10 nodes for each job using 500MB-3000MB of memory)
            - National Water Model model runs for analysis and assimilation, short-range, medium-range and long-range configurations (MPI with 32 nodes for each job, using 100 MB memory)

- Runtimes(on WCOSS)
  - USGS streamflow processing
    - JNWM_USGS_TIMESLICES: ~ 6 minutes(1 node - 24 cores per node)

  - Forcing engine
    - JNWM_FORCING_RADARONLY: ~ < 1 minute(4 nodes - 8 cores per node)
    - JNWM_FORCING_RADARGAUGE: ~ < 1 minute(4 nodes - 8 cores per node)
    - JNWM_FORCING_ANALYSIS_ASSIM: ~ 2 minutes(1 node - 18 cores per node)
    - JNWM_FORCING_SHORT_RANGE: ~ 30 minutes(3 nodes - 8 cores per node)
    - JNWM_FORCING_MEDIUM_RANGE: ~ 10 minutes(10 nodes - 12 cores per node)
    - JNWM_FORCING_LONG_RANGE (Long Range 1-4): ~ 80 minutes(8 nodes - 12 cores per node)
    - JNWM_FORCING_MPE: ~ 15 minutes (1 node - 1 core)
    - JNWM_FORCING_ANALYSIS_ASSIM_EXTEND: ~ 4 minutes (2 nodes - 14 cores per node)
    - JNWM_FORCING_HAWAII: ~ 4 minutes (3 nodes - 8 cores per node)
  - Model runs
    - JNWM_MODEL (Analysis and Assimilation): ~ 10 minutes(32 nodes - 24 cores per node)
    - JNWM_MODEL (Short Range): ~15 minutes (32 nodes - 24 cores per node)
    - JNWM_MODEL (Medium Range 1-7): ~ 75 minutes (32 nodes - 24 cores per node)
    - JNWM_MODEL (Long Range 1-4): ~ 50 minutes (32 nodes - 24 cores per node)
    - JNWM_MODEL (Extended Analysis and assimilation): ~ 25 minutes (32 nodes - 24 cores per node)
    - JNWM_MODEL (Long-range Analysis and assimilation): ~ 15 minutes (20 nodes - 24 cores per node)

                - JNWM_MODEL (Hawaii Analysis and assimilation): ~ 10 minutes (1 node - 9 cores per node)
                - JNWM_MODEL (Hawaii Short Range): ~ 35 minutes (1 node - 9 cores per node)
        - Disk space required per day
            - COM area: ~2.4 Tb/day of forcing and output files (but only ~900GB is transferred to NOMADS each day)
            - /nwges/ area for recent files: ~3.4 Tb/day (2.9Tb nwm; 578 Gb nwm_forcing)
        - Cycling frequency
            - Every 15 minutes: USGS streamflow observation processing
            - Hourly: analysis and assimilation forcing engine and model, short range forcing engine and model; Hawaii analysis and assimilation model.
            - Every 6 hours: long range analysis and assimilation; medium-range forcing engine and model; Hawaii forcing and Hawaii short range model.
            - Daily: long-range forcing and model (4 members executed four times per day for a total of 16 ensemble members valid for same time period); Extended analysis assimilation forcing and model (executed 1 time per day)
        - Version of libs/compilers/shared code being used
            - Intel Fortran 15.0.3.187
            - NetCDF 4.2
            - Python 2.7.14
            - Bash 4.2.53
            - Python netcdf4 package 1.2.2
            - Python numpy package 1.10.2
            - Python ordereddict package 1.1
            - HDF5 1.8.9
            - NCL 6.3.0
        - Data retention
            - Files in /com should be retained for 10 days
            - Files in /nwges should be retained for 5 days
- Pre-implementation testing requirements
    - The following jobs should be tested
        - USGS streamflow observation processing (JNWM_USGS_TIMESLICES)

        - Analysis and assimilation forcing engine (JNWM_FORCING_ANALYSIS_ASSIM, JNWM_FORCING_RADARGAUGE, JNWM_FORCING_RADARONLY)
        - Short-range forcing engine (JNWM_FORCING_SHORT_RANGE)
        - Medium-range forcing engine (JNWM_FORCING_MEDIUM_RANGE)
        - Long-range forcing engine (JNWM_FORCING_LONG_RANGE)
        - Extended forcing (JNWM_FORCING_MPE, JNWM_FORCING_ANALYSIS_ASSIM_EXTEND)
        - Hawaii Analysis and assimilation and Hawaii Short-range forcing (JNWM_FORCING_HAWAII)

        - Analysis and assimilation model run (JNWM_MODEL)
        - Extended Analysis and assimilation model run (JNWM_MODEL)
        - Long-range Analysis and assimilation model run (JNWM_MODEL)
        - Short-range model run (JNWM_MODEL)
        - Medium-range model run (JNWM_MODEL)
        - Long-range model run (JNWM_MODEL)
        - Hawaii Analysis and assimilation and Short-range model run (JNWM_MODEL)

There are no other changes for V2.0.0 and so the release notes for previous versions are included below for reference about the overall NWM system.

V1.2.3 - Code handoff to NCO on October 26th, 2018.

- List of enhancements
    - Fixed the race-condition issue where, in some cases, the analysis assim job can be trying to update the restart files while simultaneously the short range job is reading from the same set of restart files. The fix was accomplished by:
        - Adding a new restart directory, `restart_forecast`, under `$COM` directory. Forecast jobs will only read restart files from this directory.
        - The analysis assim job will still use the `restart` directory for reading and updating the tm02 restart files for use by the next analysis cycle. This is unchanged from the previous version of the software. The new alteration is that this same job now writes the tm00 restart files to the `restart_forecast` directory for use in initializing forecast jobs.

    - Changed the copy and link commands in `drv.sh` and `back_restart.sh` that copy or update restart files from `cp` or `ln` to the NCO copy utility, `cpfs`.
    - Note that run times will be comparable to the values cited in the job runtime section below.
    - Changed error messages in forcing analysis assim job when the data from previous cycles of forcing short range jobs is missing. The messages indicate additional details about failed copy commands and missing data in the `$NWGES` directory that can be used to diagnose issues.

There are no other changes for V1.2.3 and so the release notes for previous versions are included below for reference about the overall NWM system.

V1.2.0 - Code handoff to NCO on November 3rd, 2017 with update on January 16th, 2017. Implementation in March 2018.

- List of enhancements
    - Refined land surface and hydrologic parameters by expanding calibration from ~40 to ~1000+ calibration basins and improving parameter regionalization process.
    - Improved streamflow data assimilation in Analysis cycle
    - Enabled hooks within system structure to accept on-demand updates of hydrologic parameters
    - Various hydrofabric improvements including
        - Fixing 500 stream breaks
        - Adding 12,468 oCONUS basins (94,824 km^2)
        - Adding 680 USGS stream gauges into assimilation routine
    - Added three output variables (QSfcLatRunoff, QBucket, QBtmVertRunoff) to channel output files to support channel-only configuration of NWM used in research and development activities
    - Improved adherence to WCOSS standards via addressing NCO bugzilla tickets
    - Changed the forcing output metadata variable “time” data type from “double” to “integer”.
    - Software enhancements including
        - Streamlined compression and thinning of files
        - Moved new compression-related environment variables (COMPRESSION, COMINDPARM and GESINSPINRST) from ecf/model_envir.h to /jobs scripts.
        - Fixed USGS data processing script such that it no longer ignores secondary USGS quality flags and removed a dependency on xml.util.iso8601 that is not available in Python 2.7
- Content of output files varies by forecast configuration (internal NetCDF variable names and units are given in parenthesis)
    - Analysis and Assimilation

        - 1km gridded NetCDF
            - Near surface soil moisture deficit (40cm thickness) (soilsat_top, fraction)
            - Accumulated ET (accet, mm)
            - Snow temperature - column integrated (snowt_avg, degrees K)
            - Column averaged snow cover fraction (fsno, fraction)
            - Snow water equivalent (sneqv, kg/m2)
            - Snow depth (snowh, m)
        - 250m gridded NetCDF
            - Ponded water depth (sfcheadsubrt, mm)
            - Depth to soil saturation (zwattablrt, m)
        - Point-type channel NetCDF
            - Streamflow (streamflow, m3/sec)
            - Streamflow data assimilation increment (nudge, m3/sec)
            - Stream velocity (velocity, m/s)
            - Runoff from terrain routing (qSfcLatRunoff, m3)
            - Runoff from bottom of soil to gw bucket (qBtmVertRunoff, m3)
            - Flux from gw bucket (qBucket, m3)
        - Point-type reservoir NetCDF
            - Reservoir water surface elevation (elevation, m)
            - Reservoir inflow (inflow, m3/sec)
            - Reservoir outflow (outflow, m3/sec)
    - Short-Range
        - 1km gridded NetCDF
            - Near surface soil moisture saturation (40cm thickness) (soilsat_top, fraction)
            - Accumulated ET (accet, mm)
            - Snow temperature - column integrated (snowt_avg, degrees K)
            - Column averaged snow cover fraction (fsno, fraction)
            - Snow water equivalent (sneqv, kg/m2)
            - Snow depth (snowh, m)
        - 250m gridded NetCDF
            - Ponded water depth (sfcheadsubrt, mm)
            - Depth to soil saturation (zwattablrt, m)
        - Point-type channel NetCDF
            - Streamflow (streamflow, m3/sec)
            - Streamflow data assimilation increment (nudge, m3/sec)
            - Stream velocity (velocity, m/s)
            - Runoff from terrain routing (qSfcLatRunoff, m3)
            - Runoff from bottom of soil to gw bucket (qBtmVertRunoff, m3)
            - Flux from gw bucket (qBucket, m3)
        - Point-type reservoir NetCDF
            - Reservoir water surface elevation (elevation, m)

                - Reservoir inflow (inflow, m3/sec)
                - Reservoir outflow (outflow, m3/sec)
    - Medium-Range
        - 1km gridded NetCDF
            - Accumulated underground runoff (ugdrnoff, mm)
            - Accumulated snowmelt (acsnom, mm)
            - Snow depth (snowh, m)
            - Snow water equivalent (sneqv, kg/m2)
            - Total canopy water storage (canwat, mm)
            - Accumulated canopy evaporation (accecan, mm)
            - Accumulated transpiration (accetran, mm)
            - Accumulated direct soil evaporation (accedir, mm)
            - Snow layer liquid water (snliq, mm)
            - Number of snow layers (isnow)
            - Soil temperature (on native layers) (soil_t, degrees K)
            - Snow temperature - column integrated (snowt_avg, degrees K)
            - Snow cover fraction (fsno, fraction)
            - Volumetric soil moisture (on native layers) (soil_m, m3/m3)
            - Near surface soil moisture deficit, 40cm thickness (soilsat_top, fraction)
            - Soil ice fraction - column integrated (soilice, fraction)
            - Accumulated evapotranspiration (accet, mm)
            - Ground heat flux (grdflx, w/m2)
            - Sensible heat flux (hfx, w/m2)
            - Latent heat flux (lh, w/m2)
            - Net longwave flux (fira, w/m2)
            - Net shortwave flux (fsa, w/m2)
            - Surface radiative temperature (trad, degrees K)
        - 250m gridded NetCDF
            - Ponded water depth (sfcheadsubrt, mm)
            - Depth to soil saturation (zwattablrt, m)
        - Point-type channel NetCDF
            - Streamflow (streamflow, m3/sec)
            - Streamflow data assimilation increment (nudge, m3/sec)
            - Stream velocity (velocity, m/s)
        - Point-type reservoir NetCDF
            - Reservoir water surface elevation (elevation, m)
            - Reservoir inflow (inflow, m3/sec)
            - Reservoir outflow (outflow, m3/sec)
    - Long-Range
        - 1km gridded NetCDF
            - Underground runoff (ugdrnoff, mm)
            - Surface Runoff (sfcrnoff, mm)

                - Accumulated snowmelt (acsnom, mm)
                - Snow water equivalent (sneqv, kg/m2)
                - Total column soil moisture deficit (soilsat, fraction)
                - Total canopy water storage (canwat, mm)
                - Near surface soil moisture deficit, 40cm thickness (soilsat_top, fraction)
                - Accumulated evapotranspiration (accet, mm)
            - 250m gridded NetCDF
                - None
            - Point-type channel NetCDF
                - Streamflow (streamflow, m3/sec)
                - Streamflow data assimilation increment (nudge, m3/sec)
                - Stream velocity (velocity, m/s)
            - Point-type reservoir NetCDF
                - Reservoir water surface elevation (elevation, m)
                - Reservoir inflow (inflow, m3/sec)
                - Reservoir outflow (outflow, m3/sec)
    - All of the preceding analysis and forecast configurations require the same set of 1km forcing variables (precipitation rate, surface pressure, shortwave radiation, longwave radiation, u-wind, v-wind, temperature, and specific humidity) but, as outlined above, these fields come from different sources for each model configuration. A selection of forcing files are disseminated alongside the model analysis and short/medium forecast output.
        - 2 meter temperature (t2d, K)
        - 2 meter specific humidity (q2d, kg/kg)
        - 10 meter U wind component (u2d, m/sec)
        - 10 meter V wind component (v2d, m/sec)
        - Surface pressure (psfc, Pa)
        - Downward shortwave radiation (swdown, w/m2)
        - Downward longwave radiation (lwdown, w/m2)
        - Precipitation rate (rainrate, kg/m2sec)
- Compute resource information
    - Job type (serial or parallel) and associated node use
        - Serial
            - USGS streamflow observation processing (1 node)
        - Parallel
            - Forcing engine jobs for analysis and assimilation, short-range, medium-range and long-range configurations (1-10 nodes for each job using 500MB-3000MB of memory)
            - National Water Model model runs for analysis and assimilation, short-range, medium-range and long-range configurations (MPI with 32 nodes for each job, using 100 MB memory)
    - Runtimes(on WCOSS)

        - USGS streamflow processing
            - JNWM_USGS_TIMESLICES: ~ 22 minutes(1 node - 24 cores per node)
        - Forcing engine
            - JNWM_FORCING_RADARONLY: < 4 minute(4 nodes - 8 cores per node)
            - JNWM_FORCING_RADARGAUGE: < 3 minute(4 nodes - 8 cores per node)
            - JNWM_FORCING_ANALYSIS_ASSIM: < 4 minutes(1 node - 18 cores per node)
            - JNWM_FORCING_SHORT_RANGE: ~ 36 minutes(3 nodes - 8 cores per node)
            - JNWM_FORCING_MEDIUM_RANGE: ~ 16 minutes(10 nodes - 12 cores per node)
            - JNWM_FORCING_LONG_RANGE (Long Range 1-4): ~ 82 minutes(8 nodes - 12 cores per node)
        - Model runs
            - JNWM_MODEL (Analysis and Assimilation): ~21 minutes(32 nodes - 24 cores per node)
            - JNWM_MODEL (Short Range): ~31 minutes (32 nodes - 24 cores per node)
            - JNWM_MODEL (Medium Range): ~ 100 minutes (32 nodes - 24 cores per node)
            - JNWM_MODEL (Long Range 1-4): ~ 57 minutes (32 nodes - 24 cores per node)
    - Disk space required per day
        - COM area for transfer to NOMADS: ~387 Gb/day of forcing and output files
        - /nwges/ area for recent files: ~1.3 Tb/day
    - Cycling frequency
        - Every 15 minutes: USGS streamflow observation processing
        - Hourly: analysis and assimilation forcing engine and model, short range forcing engine and model
        - Every 6 hours: medium-range forcing engine and model
        - Daily: long-range forcing and model (4 members executed four times per day for a total of 16 ensemble members valid for same time period)
    - Version of libs/compilers/shared code being used
        - Intel Fortran 15.0.3.187
        - NetCDF 4.2
        - Python 2.7.12
        - Bash 4.2.53
        - Python netcdf4 package 1.2.2
        - Python numpy package 1.10.2

        - Python ordereddict package 1.1
        - HDF5 1.8.9
        - NCL 6.3.0
    - Data retention
        - Files in /com should be retained for 10 days
        - Files in /nwges should be retained for 5 days
- Pre-implementation testing requirements
    - The following jobs should be tested
        - USGS streamflow observation processing (JNWM_USGS_TIMESLICES)
        - Analysis and assimilation forcing engine (JNWM_FORCING_ANALYSIS_ASSIM, JNWM_FORCING_RADARGAUGE, JNWM_FORCING_RADARONLY)
        - Short-range forcing engine (JNWM_FORCING_SHORT_RANGE)
        - Medium-range forcing engine (JNWM_FORCING_MEDIUM_RANGE)
        - Long-range forcing engine (JNWM_FORCING_LONG_RANGE)
        - Analysis and assimilation model run (JNWM_MODEL)
        - Short-range model run (JNWM_MODEL)
        - Medium-range model run (JNWM_MODEL)
        - Long-range model run (JNWM_MODEL)

There are no other changes for V1.2.0 and so the release notes for previous versions are included below for reference about the overall NWM system.

V1.1.4 - released July 27th, 2017

- Fix to allow model to ingest USGS stream gauge observation files from previous day. Necessary for Analysis execution near 00Z day boundary.
- Update to crontab file to match execution times now used by NCO in operations. Crontab used if non-ECFlow (cronjob-based) execution of system is needed.

V1.1.3 - released March 3rd, 2017

- ix for 00z model analysis failure associated with missing previous day forcing data
- Only copy or alert USGS timeslice files that are newer than 4 hours and not being modified by the current job to the COM dir
- Fix for bug which led to corrupt data along the HRRR/RAP boundary after combining

V1.1.2 - released February 22nd, 2017

- Corrected the internal time stamps inside of the forcing Analysis and Assimilation files
- Changed NWM_NCDF_TGZ to NWM_NCDF

V1.1.1 - released January 25th, 2017

- Increased USGS job run time limit to 40 minutes
- Moved reforecast job and exec scripts to ecf/testReforecast directory

- Changed CASETYPE, parm subdirectories and output directories under com/ and nwges/ to consistent with CASETYPE (now all in lowercase).

v1.1.0 - released January 9th, 2017

- First full upgrade of the National Water Model
- Will enable NWS to provide improved high resolution, seamless nationwide analyses and forecasts of streamflow and other water resource variables

List of Enhancements:

- Forecast cycling
    - Extension of short range forecast from 15 to 18 hours
    - Increased frequency of medium range forecast from 1 to 4 times per day
- Parameter Updates
    - Parameter calibration to reduce hydrologic biases
    - Improvement of MRMS/HRRR/RAP precipitation blending in NWM Analysis
    - Improvement of HRRR/RAP precipitation blend in NWM Short Range forecast
    - Incorporated a data-driven approach to channel parameter estimates utilizing a regression model based on USGS cross-sections.
- Stream Connectivity Refinements
    - Corrections to NHDPlusV2 streamflow connectivity errors noted by RFCs
    - Addition of new oCONUS hydrologically contributing areas
- Additional upgrades
    - Fix for overly large infiltration rate in sandy areas during intense precipitation
    - Fixes to underlying topography leading to isolated extreme OCONUS ponded water values
    - Fix for high altitude excessive snow melt
    - Fix for isolated negative underground runoff
    - Header and NetCDF format updates to improve compatibility with NetCDF file readers
    - Enhancements to USGS data decoder and preprocessor
    - Use of scaling and offset factors along with conversion of output files to NetCDF4 to reduce file sizes and increase accessibility
    - Improved adherence to WCOSS standards via addressing NCO bugzilla tickets
    - Corrected locations and stream reach attributions to several previously Compute resource information
        - Job type (serial or parallel) and associated node use
            - Serial
                - USGS streamflow observation processing (1 core)
            - Parallel
                - Forcing engine jobs for analysis and assimilation, short-range, medium-range and long-range configurations

(1-10 nodes for each job using 500MB-3000MB of memory)
        - National Water Model model runs for analysis and assimilation, short-range, medium-range and long-range configurations (MPI with 32 nodes for each job, using 100 MB memory)
- Runtimes
    - USGS streamflow processing
        - JNWM_USGS_TIMESLICES: ~ 23 minutes
    - Forcing engine
        - JNWM_FORCING_RADARONLY: < 2 minute
        - JNWM_FORCING_RADARGAUGE: < 2 minute
        - JNWM_FORCING_ANALYSIS_ASSIM: < 3 minutes
        - JNWM_FORCING_SHORT_RANGE: ~ 35 minutes
        - JNWM_FORCING_MEDIUM_RANGE: ~ 23 minutes
        - JNWM_FORCING_LONG_RANGE (Long Range 1-4): ~ 95 minutes
    - Model runs
        - JNWM_MODEL (Analysis and Assimilation): ~22 minutes
        - JNWM_MODEL (Short Range): ~35 minutes
        - JNWM_MODEL (Medium Range): ~ 105 minutes
        - JNWM_MODEL (Long Range 1-4): ~ 63 minutes
- Disk space required per day
    - COM area for transfer to NOMADS: ~379 Gb/day of forcing and output files
    - /nwges/ area for recent files: ~3.3 Tb/day
- Cycling frequency
    - Every 15 minutes: USGS streamflow observation processing
    - Hourly: analysis and assimilation forcing engine and model, short range forcing engine and model
    - Every 6 hours: medium-range forcing engine and model
    - Daily: long-range forcing and model (4 members executed four times per day for a total of 16 ensemble members valid for same time period)
- Version of libs/compilers/shared code being used
    - Intel Fortran 15.0.3.187
    - NetCDF 4.2
    - Python 2.6.9
    - Bash 4.3.53
    - Python netcdf4 package 1.2.2
    - Python numpy package 1.10.2
    - Python ordereddict package 1.1
    - HDF5 1.8.9

            - NCL 6.3.0
        - Data retention
            - Files in /com should be retained for 10 days
            - Files in /nwges should be retained for 1 day
    - Pre-implementation testing requirements
        - The following jobs should be tested
            - USGS streamflow observation processing (JNWM_USGS_TIMESLICES)
            - Analysis and assimilation forcing engine (JNWM_FORCING_ANALYSIS_ASSIM, JNWM_FORCING_RADARGAUGE, JNWM_FORCING_RADARONLY)
            - Short-range forcing engine (JNWM_FORCING_SHORT_RANGE)
            - Medium-range forcing engine (JNWM_FORCING_MEDIUM_RANGE)
            - Long-range forcing engine (JNWM_FORCING_LONG_RANGE)
            - Analysis and assimilation model run (JNWM_MODEL)
            - Short-range model run (JNWM_MODEL)
            - Medium-range model run (JNWM_MODEL)
            - Long-range model run (JNWM_MODEL)
- Dissemination information
    - One major change from V1.0.6: Files are now provided in the NetCDF4 format, with contents scaled and compressed. This eliminates the need for gzipping and greatly reduces file size. Files feature improved compatibility with graphical display tools like ncview, Panoply and Weather and Climate Toolkit, along with GIS tools.

There are no other changes for V1.1.0 and so the release notes for previous versions are included below for reference about the overall NWM system.

v1.0.6 - released December 9th, 2016

- Alterations to USGS ingest scripts supporting National Water Model usage of USGS streamflow observations
- Necessitated by upcoming changes to USGS XML data format, allows NWM to continue to parse the modified XML format to retrieve the needed streamflow observations and continue current operations. I.e., not an enhancement, but rather an updated necessitated by an upstream format change.
- Changes are isolated to the USGS-related scripts. All file names remain the same, and files have not been added or removed--only the contents of the files have changed.

There are no other changes for V1.0.6 and so the release notes for V1.0 are included below for reference about the overall NWM system.

v1.0 - released June 14, 2016

- Initial implementation of the National Water Model
- Will enable NWS to provide experimental high resolution, seamless nationwide analyses and forecasts of streamflow for the first time
- Streamflow analyses and forecasts will be produced for 2.7 million stream reaches, and other hydrologic information across the CONUS will be output on 1km and 250m grids.
- Uses the community WRF-Hydro model developed by NCAR as core modeling architecture.
- WRF-Hydro is configured to use the NoahMP LSM on a 1km grid, diffusive overland and saturated subsurface routing on a 250m grid, and Muskingum-Cunge channel routing on NHDPlusV2 stream reaches. 1,260 reservoirs are represented in the model.
- The NWM will execute in four configurations
    - Analysis and Assimilation
        - Cycles hourly and outputs a one-hour snapshot of current hydrologic conditions
        - Forced with MRMS gauge-adjusted and radar-only QPE along with HRRR and RAP short-range forecasts for filling gaps in MRMS coverage
        - Assimilates USGS streamflow observations from ~8,000 real-time gauges. Data are obtained from the USGS and are considered provisional at the time they were collected. USGS observations are subject to change as new data and field observations are made and incorporated into the USGS data computations.
        - Produces a single restart file used in common by all NWM forecast simulations
    - Short-Range Forecast
        - Cycles hourly and produces hourly hydrologic deterministic forecasts out to 15 hours
        - Forced with HRRR and RAP NWP model forecasts
    - Medium-Range Forecast
        - Cycles daily and produces 3-hourly deterministic hydrologic forecasts out to 10 days
        - Forced with GFS NWP model forecasts
    - Long-Range Forecast
        - Cycles four times per day (each time with four members) and produces 6-hourly ensemble hydrologic forecasts out to 30-days
        - Forced with the 16 CFS forecasts produced each day by NCEP
- Output Overview
    - All output will be produced in NetCDF format in one of four formats
        - 1km gridded NetCDF covering a CONUS+ domain that extends roughly north to latitude 58N and south to latitude 18N.
        - 250m gridded NetCDF covering a CONUS+ domain that extends rough north to latitude 58N and south to latitude 18N.

        - Point-type NetCDF for 2.7 million NHDPlusV2 stream reaches covering the CONUS and hydrologically contributing areas
        - Point-type NetCDF for 1260 reservoirs that have connectivity to the NHDPlusV2 channel network
    - Content of output files varies by forecast configuration (internal NetCDF variable names and units are given in parenthesis)
        - Analysis and Assimilation
            - 1km gridded NetCDF
                - Near surface soil moisture deficit (40cm thickness) (soilsat_top, fraction)
                - Accumulated ET (accet, mm)
                - Snow temperature - column integrated (snowt_avg, degrees K)
                - Column averaged snow cover fraction (fsno, fraction)
                - Snow water equivalent (sneqv, kg/m2)
                - Snow depth (snowh, m)
            - 250m gridded NetCDF
                - Ponded water depth (sfcheadsubrt, mm)
                - Depth to soil saturation (zwattablrt, m)
            - Point-type channel NetCDF
                - Streamflow (streamflow, m3/sec)
                - Streamflow data assimilation increment (nudge, m3/sec)
                - Stream velocity (velocity, m/s)
                - Channel inflow (q_lateral, m3/sec)
            - Point-type reservoir NetCDF
                - Reservoir water surface elevation (elevation, m)
                - Reservoir inflow (inflow, m3/sec)
                - Reservoir outflow (outflow, m3/sec)
        - Short-Range
            - 1km gridded NetCDF
                - Near surface soil moisture deficit (40cm thickness) (soilsat_top, fraction)
                - Accumulated ET (accet, mm)
                - Snow temperature - column integrated (snowt_avg, degrees K)
                - Column averaged snow cover fraction (fsno, fraction)
                - Snow water equivalent (sneqv, kg/m2)
                - Snow depth (snowh, m)
            - 250m gridded NetCDF
                - Ponded water depth (sfcheadsubrt, mm)
                - Depth to soil saturation (zwattablrt, m)
            - Point-type channel NetCDF
                - Streamflow (streamflow, m3/sec)

        - Streamflow data assimilation increment (nudge, m3/sec)
        - Stream velocity (velocity, m/s)
        - Channel inflow (q_lateral, m3/sec)
    - Point-type reservoir NetCDF
        - Reservoir water surface elevation (elevation, m)
        - Reservoir inflow (inflow, m3/sec)
        - Reservoir outflow (outflow, m3/sec)
- Medium-Range
    - 1km gridded NetCDF
        - Accumulated underground runoff (ugdrnoff, mm)
        - Accumulated snowmelt (acsnom, mm)
        - Snow depth (snowh, m)
        - Snow water equivalent (sneqv, kg/m2)
        - Total canopy water storage (canwat, mm)
        - Accumulated canopy evaporation (accecan, mm)
        - Accumulated transpiration (accetran, mm)
        - Accumulated direct soil evaporation (accedir, mm)
        - Snow layer liquid water (snliq, mm)
        - Number of snow layers (isnow)
        - Soil temperature (on native layers) (soil_t, degrees K)
        - Snow temperature - column integrated (snowt_avg, degrees K)
        - Snow cover fraction (fsno, fraction)
        - Volumetric soil moisture (on native layers) (soil_m, m3/m3)
        - Near surface soil moisture deficit, 40cm thickness (soilsat_top, fraction)
        - Soil ice fraction - column integrated (soilice, fraction)
        - Accumulated evapotranspiration (accet, mm)
        - Ground heat flux (grdflx, w/m2)
        - Sensible heat flux (hfx, w/m2)
        - Latent heat flux (lh, w/m2)
        - Net longwave flux (fira, w/m2)
        - Net shortwave flux (fsa, w/m2)
        - Surface radiative temperature (trad, degrees K)
    - 250m gridded NetCDF
        - Ponded water depth (sfcheadsubrt, mm)
        - Depth to soil saturation (zwattablrt, m)
    - Point-type channel NetCDF
        - Streamflow (streamflow, m3/sec)
        - Streamflow data assimilation increment (nudge, m3/sec)
        - Stream velocity (velocity, m/s)
        - Channel inflow (q_lateral, m3/sec)
    - Point-type reservoir NetCDF

        - Reservoir water surface elevation (elevation, m)
        - Reservoir inflow (inflow, m3/sec)
        - Reservoir outflow (outflow, m3/sec)
- Long-Range
    - 1km gridded NetCDF
        - Underground runoff (ugdrnoff, mm)
        - Surface Runoff (sfcrnoff, mm)
        - Accumulated snowmelt (acsnom, mm)
        - Snow water equivalent (sneqv, kg/m2)
        - Total column soil moisture deficit (soilsat, fraction)
        - Total canopy water storage (canwat, mm)
        - Near surface soil moisture deficit, 40cm thickness (soilsat_top, fraction)
        - Accumulated evapotranspiration (accet, mm)
    - 250m gridded NetCDF
        - None
    - Point-type channel NetCDF
        - Streamflow (streamflow, m3/sec)
        - Streamflow data assimilation increment (nudge, m3/sec)
        - Channel inflow (q_lateral, m3/sec)
    - Point-type reservoir NetCDF
        - Reservoir water surface elevation (elevation, m)
        - Reservoir inflow (inflow, m3/sec)
        - Reservoir outflow (outflow, m3/sec)
- Note that in areas with very little snow cover, the snow temperature will represent a blend between the snow and surface temperature. In isolated cases, this may result in combined temperature values that exceed 273.16 deg. K.
- All of the preceding analysis and forecast configurations require the same set of 1km forcing variables (precipitation rate, surface pressure, shortwave radiation, longwave radiation, u-wind, v-wind, temperature, and specific humidity) but, as outlined above, these fields come from different sources for each model configuration. A selection of forcing files are disseminated alongside the model analysis and short/medium forecast output.
        - 2 meter temperature (t2d, K)
        - 2 meter specific humidity (q2d, kg/kg)
        - 10 meter U wind component (u2d, m/sec)
        - 10 meter V wind component (v2d, m/sec)
        - Surface pressure (psfc, Pa)
        - Downward shortwave radiation (swdown, w/m2)
        - Downward longwave radiation (lwdown, w/m2)
        - Precipitation rate (rainrate, kg/m2sec)

- Compute resource information
    - Job type (serial or parallel) and associated node use
        - Serial
            - USGS streamflow observation processing (1 core)
        - Parallel
            - Forcing engine jobs for analysis and assimilation, short-range, medium-range and long-range configurations (4-10 nodes for each job using 500MB-3000MB of memory)
            - National Water Model model runs for analysis and assimilation, short-range, medium-range and long-range configurations (MPI with 32 nodes for each job, using 100 MB memory)
    - Runtimes
        - USGS streamflow processing
            - JNWM_USGS_TIMESLICES: ~21 minutes
        - Forcing engine
            - JNWM_FE_RADARONLY: <1 minute
            - JNWM_FE_RADARGAUGE: <1 minute
            - JNWM_FE_QPEBIAS: <1 minutes
            - JNWM_FE_RAPHRRR: ~35 minutes
            - JNWM_FE_GFS: ~11 minutes
            - JNWM_FE_CFS (Long Range 1-4): ~95 minutes
        - Model runs
            - JNWM_MODEL (Analysis and Assimilation): ~24 minutes
            - JNWM_MODEL (Short Range): ~40 minutes
            - JNWM_MODEL (Medium Range): ~92 minutes
            - JNWM_MODEL (Long Range 1-4): ~56 minutes
    - Disk space required per day
        - COM area for transfer to NOMADS: ~600 Gb/day of forcing and output files
        - /nwges/ area for recent files: ~2.8 Tb/day
    - Cycling frequency
        - Every 15 minutes: USGS streamflow observation processing
        - Hourly: analysis and assimilation forcing engine and model, short range forcing engine and model
        - Daily: medium-range forcing engine and model, long-range forcing and model (4 members executed four times per day for a total of 16 ensemble members valid for same time period)
    - Version of libs/compilers/shared code being used
        - Intel Fortran 15.0.3.187
        - NetCDF 4.2
        - Python 2.7.12
        - Bash 4.3.53

        - Python netcdf4 package 1.2.2
        - Python numpy package 1.10.1
        - Python ordereddict package 1.1
        - HDF5 1.8.9
        - NCL 6.3.0
    - Data retention
        - Files in /com should be retained for 10 days
        - Files in /nwges should be retained for 1 day
- Pre-implementation testing requirements
    - The following jobs should be tested
        - USGS streamflow observation processing (JNWM_USGS_TIMESLICES)
        - Analysis and assimilation forcing engine (JNWM_FE_QPEBIAS, JNWM_FE_RADARGAUGE, JNWM_FE_RADARONLY)
        - Short-range forcing engine (JNWM_FE_RAPHRRR)
        - Medium-range forcing engine (JNWM_FE_GFS)
        - Long-range forcing engine (JNWM_FE_CFS)
        - Analysis and assimilation model run (JNWM_MODEL)
        - Short-range model run (JNWM_MODEL)
        - Medium-range model run (JNWM_MODEL)
        - Long-range model run (JNWM_MODEL)
        - This is an initial implementation, and so does require a 30-day evaluation
    - Suggested evaluators include: RFCs, WFOs, NCEP/EMC, and NCEP/WPC
- Dissemination information
    - Output
        - The full selection of model forcing and model output files will be sent to NOMADS and the NCEP ftp server. Some files will be gzipped as needed.
        - Due to storage space limitations, the latitude and longitude of each output grid-box (for grid-type data) and point (for point-type data) are stored outside of the individual NWM NOMADS NetCDF files, in files available at ftp://ftp.nohrsc.noaa.gov/pub/staff/keicher/WRFH_ppd/web/NWM_nc_tools.tar.gz . Scripts are available at this same location which will append this geospatial data to a user selected output file, enabling viewing within common NetCDF visualization utilities.
        - A limited selection of files will be sent to the HPSS tape system (see below)
        - There will be three main routes for information dissemination
            - NOMADS file server
            - NWC website map viewer at http://water.noaa.gov/map (will draw data from NOMADS)
            - RFC CHPS systems (will utilize thinned NWC data feed which in turn obtains full data from NOMADS)

        - Day 1 users include NWS RFC and WFO field offices, NWS regions and NCEP centers. Follow-on users include the USGS, FEMA, EPA, ACE, NASA, USBR, state agencies, local and regional emergency responders, the National Ocean Service, NIDIS, academia and private industry.
    - The files that need to be transferred from prod WCOSS to dev WCOSS to support timely model restarts upon a WCOSS switch include: model LSM restart files, model routing restart files, processed USGS streamflow observations, 1km forcing files for each of the four configurations described above.
    - Code and data status: The code is not proprietary and no restricted data is used
    - Log files will be written out to the /com directory structure for use by the NWC, but not publicly disseminated.
- Archive to HPSS
    - A limited set of data will be stored within a 1-year rotating archive on HPSS using the tape space approved by the HPCRAC.
        - Channel routing and LSM restart files
        - Forcing data from all configurations
            - Full forcing files from analysis and assimilation, short-range and medium-range configurations
            - Thinned forcing files (temperature and precipitation) from all 16 CFS forcing members
        - Processed USGS streamflow observations
        - Streamflow, stream velocity, nudging increment and channel inflow
        - Ponded water, depth to saturation