HEFS

Overview and Getting Started Manual

Version: HEFS 1.2.1

Release Date: 28 October 2014

National Weather Service
Office of Hydrologic Development

Table of Contents

1 Overview 3

1.1 Notation 3

1.2 Terminology 3

2 Introduction 4

3 What are the HEFS Software Components? 7

3.1 MEFP 7

3.1.1 Description 7
3.1.2 MEFPPE 8
3.1.3 MEFP Data Ingest Components 10
3.1.4 MEFP Forecast Components 12

3.2 EnsPost 14

3.2.1 Description 14
3.2.2 EnsPostPE 14
3.2.3 EnsPost Forecast Components 17

3.3 EVS 17

3.3.1 Description 17
3.3.2 EVS-to-CHPS model adapter 19
3.3.3 Hindcasting 19

4 What Next? 22

4.1 HEFS System 22

4.2 Installing HEFS 22

Appendix A: Limitations 24

1 Overview

This guide provides an introduction to the Hydrologic Ensemble Forecast Service (HEFS) suite of software, describing in general terms the components of HEFS and their functions. A detailed presentation of each component is provided in the *User's Manual* associated with that component and delivered as part of the release of HEFS. Instructions for configuring the components to be a part of CHPS is provided in the *Configuration Guide*(s) associated with that component and delivered as part of the release. All users should read Sections 2, 3, and 4 before installing the HEFS release and configuring and using any HEFS software component.

1.1 *Notation*

Within this document, the following notation is used:

- All graphical interface components are **Capitalized and in Bold**.
- All important terms are *italicized* when first mentioned.

1.2 *Terminology*

The following terms are used in this document:

- *installation segment*: The id of the first segment for which HEFS components will be configured.
- *installation catchments*: The catchments associated with the *installation segment*.
- *installation standalone*: The standalone used to configure MEFP and EnsPost and confirm the configuration.
- *parameter estimation standalone*: The standalone used for execution of the MEFPPE and EnsPostPE.

2 Introduction

The HEFS was developed by the Office of Hydrologic Development (OHD) of the U.S. National Weather Service (NWS). The HEFS issues hydrologic forecasts that are "uncertainty aware", i.e. they provide information about forecast uncertainty. This is achieved by issuing an ensemble of possible values of the forecast variables. Unlike single-valued or "deterministic" forecasts, which comprise a single estimate of the forecast variable at each time and location, an ensemble forecast provides a set of possible values. An ensemble forecasting system, such as the HEFS, translates or "propagates" an ensemble of inputs (e.g. precipitation and temperature) through a hydrologic model to provide an ensemble of outputs (streamflow). Ensemble forecasting provides a convenient way to quantify and trace the movement of uncertainty through hydrologic models, which otherwise require fixed values of the inputs (i.e. fixed values of temperature and precipitation). In ensemble forecasting, the hydrologic models, such as SAC-SMA and SNOW-17, are executed repeatedly. Each execution uses a different value for the inputs and, by implication, provides one possible value for the outputs (streamflow). By collecting together the ensemble of outputs, the forecasts can be used to develop probability statements, such as the probability of flooding (fraction of members that exceed flood stage), or to use in other modeling tools or decision support systems.

Ensemble forecasting relies on a combination of physically-based and statistical modeling. The HEFS comprises both physically-based hydrologic models (e.g. SAC-SMA, SNOW-17) and statistical modeling of the forecast uncertainties. The components of the HEFS are implemented within a modular software framework (Figure 1). The HEFS modules aim to quantify and, where possible, reduce the uncertainties at various stages in the hydrologic modeling process, as well as generate outputs for operational forecasting.

Statistical models rely on historical observations to determine historical forecast errors. This requires statistical modeling of the relationship between the past forecasts and observations. If this relationship is relatively constant or "stationary" in time, past forecasting errors provide a statistical guide to future forecasting errors. These statistical models provide the basis to generate the ensemble of inputs required by the hydrologic models (the Meteorological Ensemble Forecast Processor or MEFP) and to correct for consistent errors (biases) in the streamflow predictions (the Ensemble Post-processor or EnsPost). Of course, the ensembles generated by these models must be physically, as well as statistically, reasonable. In particular, they must reproduce observed patterns of forcing and streamflow in space and time. For example, adjacent basins could have similar precipitation amounts at any given time. They must also reproduce the observed relationships between variables. For example, precipitation will not fall as snow at high air temperatures. Moreover, they aim to reproduce these patterns at several different accumulation volumes (e.g. daily, monthly, etc.).

The statistical modeling in the HEFS is conducted in two parts. First, a Parameter Estimator (PE) is used to estimate the parameters of each statistical model. The parameters must be estimated from a long and consistent record of paired predictions and observations. This is necessary to minimize sampling uncertainty; that is, to provide parameter estimates that are reasonable and not too "noisy." Secondly, the estimated parameters are applied in real time to the "raw" operational forecasts, whether from the forcing models (e.g. GFS, CFSv2) or streamflow models

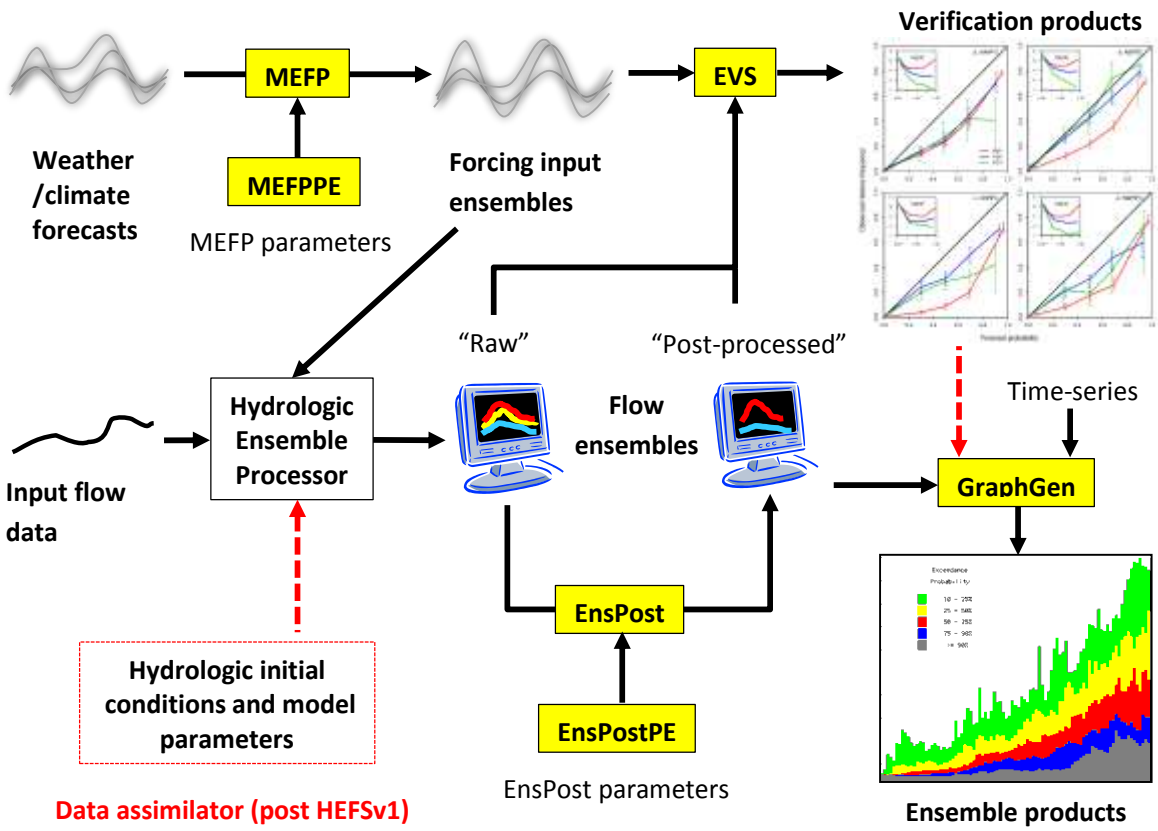

Figure 1: Overview of HEFS.

(SAC-SMA and SNOW-17). For example, the MEFP Parameter Estimator (MEFPPE) estimates the parameters of the relationship between the historical observations and the "raw" temperature and precipitation forecasts from weather and climate models, such as the Global Ensemble Forecast System (GEFS). In real time, the MEFP uses the parameters from the MEFPPE together with the raw operational forecasts from the GEFS. Using this information, the MEFP produces an ensemble of temperature and precipitation forecasts for input to the hydrologic models. The MEFP accounts for any biases in the raw forecasts via the model parameters estimated with the MEFPPE. The MEFP accommodates several sources of raw forecasts, including the single-valued operational forecasts from the RFCs, the GFS, the Climate Forecast System (v.2), historical observations ("climatological forcing") and, ~~shortly,~~ the Global Ensemble Forecast System (GEFS). These raw forecasts are integrated seamlessly into the MEFP, in order to provide bias-corrected forcing from less than 1 day out to almost one year.

For streamflow, the EnsPost Parameter Estimator (EnsPostPE) models the statistical relationship between the streamflow predictions (hydrologic simulations) and observations. Using the parameters estimated by the EnsPostPE, the EnsPost makes an adjustment to the raw streamflow forecasts in real time. This adjustment accounts for any biases identified in the historical raw forecasts used to calibrate the EnsPostPE. Both conceptually and practically, the MEFP and EnsPost are similar; they both aim to produce a forecast ensemble that is statistically consistent with the past observations under similar conditions to the "raw" inputs. However, they use

different sources of information (forcing versus streamflow) and hence account for different sources of uncertainty. In this context, the total uncertainty in hydrologic forecasting may be factored into two main sources of uncertainty and bias, namely the input or "forcing uncertainties" and the "hydrologic uncertainties." The latter comprise all sources of uncertainty and bias in the hydrologic modeling. Since the MEFP only accounts for the forcing uncertainty, the EnsPost must account for the hydrologic uncertainty. In order to model the hydrologic uncertainties separately from the forcing uncertainties, the EnsPost uses hydrologic simulations. Hydrologic simulations comprise observed forcing and hence the forcing uncertainties are effectively eliminated (or at least minimized, if the observations are relatively error-free).

As indicated above, the MEFPPE and the EnsPostPE are similar in practice, as well as conceptually. Thus, both PEs use a simplified form of statistical modeling that invokes rather strict assumptions about the processes that generated the sample data. Since both PEs model the "scatter" between two variables (observed versus forecast precipitation for MEFPPE and observed versus simulated streamflow for EnsPostPE), they model a "bivariate" relationship. Both PEs assume that this bivariate relationship follows a normal or "Gaussian" probability distribution, and model the sample data accordingly. The main advantage of this assumption is that the normal distribution has very few parameters to estimate. To assist with this, the raw data are transformed using a "Normal Quantile Transform." In other words, the statistical modeling is conducted in a space that is consistent with the normal distribution. Once the model parameters have been estimated by the PEs in transformed space, they are applied in real time (by the MEFP or the EnsPost) to generate ensemble members. Finally, the ensemble members are back-transformed into original space (e.g. to streamflow in cubic feet per second). For the MEFPPE, precipitation creates an additional complexity of being a "mixed" variable (i.e. if precipitation occurs, then it occurs with a given amount), but that is beyond the scope of this overview.

As indicated above, the HEFS aims to produce ensembles that are physically, as well as statistically, reasonable. Thus, for all temporal and spatial scales of interest, the ensemble forecasts should capture similar patterns in space and time as those observed under the same conditions. Within the MEFP, the spatial and temporal patterns in temperature and precipitation are preserved by "shuffling" the ensemble members using historical observations. This ensures the relative ordering (ranking) of the members at adjacent times or locations is consistent with those of the observations on the same dates in the historical record. The re-ordering technique is known as the "Schaake Shuffle." In the EnsPost, the temporal correlations between streamflow amounts at adjacent times are modeled with an "autoregressive" model; this exploits the persistence in streamflow over time.

Alongside the components of the HEFS that produce operational ensemble forecasts, instructions are provided for hindcasting in CHPS and a tool for verification is provided . Hindcasting is necessary to produce the long and consistent record of historical forecasts needed to properly evaluate the HEFS at particular locations. Without a long record of forecasts from a frozen version of the HEFS, it would be difficult to evaluate forecast quality with reasonably small sampling uncertainty, i.e. with reasonable confidence in the verification results. The Ensemble Verification System (EVS) allows for the verification of the HEFS hindcasts from which guidance can be developed for operational use of the HEFS (e.g. about the conditions under which performance might be impaired and what to look for).

3 What are the HEFS Software Components?

Refer to Figure 1. The HEFS software components are as follows:

- The Meteorological Ensemble Forecast Processor (MEFP): Generates forecast ensembles of precipitation and temperature for use as input for streamflow forecasting.
  - MEFP Parameter Estimator (MEFPPE): A tool to guide the users through the process of estimating parameters for the MEFP.
  - Data Ingest Components: FEWS transformations and OHD model adapter modules used to ingest operational gridded forecasts that are input to MEFP.
  - Forecast Components: FEWS transformations and OHD model adapters modules used to prepare input for MEFP, execute MEFP, and convert the generated forecast ensemble to the appropriate data types.

- The HEFS Ensemble Post-Processor (EnsPost): Post-processes an ensemble of simulated streamflows to account for hydrologic uncertainty.
  - EnsPost Parameter Estimator (EnsPostPE): A tool to guide the users through the process of estimating parameters for EnsPost.

- Ensemble Verification System: Allows for verifying the forecast outputs at various points in the HEFS process, including verifying MEFP precipitation and temperature forecast ensembles, raw streamflow forecast ensembles, and EnsPost post-processed streamflow forecast ensembles. The EVS relies on a large sample of forecasts and corresponding observations; i.e. it relies on hindcasting and/or achiving.

- CHPS Graphics Generator (GraphGen): Generate basic products summarizing inputs and/or outputs to the HEFS components. The Graphics Generator is delivered with the baseline CHPS release and will not be described further here.

Each of these components is described in greater detail in the sections that follow.

3.1 *MEFP*

3.1.1 Description

The MEFP is a statistical model that combines information from different forecast sources and generates a single forecast ensemble of meteorological inputs to streamflow models. It allows for the following three forecast sources to be used:

- 6-hour single-valued (i.e., comprised of one time series) forecasts generated by the NWS River Forecast Centers (RFC) from Hydrometeorological Prediction Center (HPC) guidance. Includes forecast of 6-hour accumulated mean areal precipitation (MAP) and 6-hour average mean areal temperature (MAT) data types.
- Gridded ensemble forecasts generated by the Global Ensemble Forecast System (GEFS) developed at the National Centers for Environmental Prediction (NCEP). This includes

both the 1998-frozen version (denoted GFS herein; forecasts of 12-hour MAP and 12-hour *instantaneous* mean areal temperature) and the recently upgraded operational version (denoted GEFS; forecasts of 6-hour MAP and 6-hour minimum and maximum temperature (TMIN and TMAX, respectively)).
- Gridded single-valued forecasts generated by the Climate Forecast System version 2 (CFSv2) of NCEP (forecasts of 6-hour MAP, TMIN, TMAX).

The forecasts produced by the forecast sources are converted to the appropriate data type:

- *precipitation*: 6-hour forecasts of MAP (FMAP)
- *temperature*: 24-hour forecasts of TMIN (TFMN) and TMAX (TFMX)

Using those time series as input, MEFP generates the following output:

- *precipitation*: A forecast ensemble of 6-hour FMAP time series.
- *temperature*: A forecast ensemble of 24-hour TFMN/TFMX time series.

The 24-hour TFMN/TFMX time series are later converted to 6-hour mean temperature (FMAT) time series to use as input to the streamflow forecasting process.

The MEFP can be used to produce ensemble forecasts from any combination of the RFC, GEFS, and CFSv2 forecast sources. In addition, it can be used to generate climatology ensemble traces with either historical observations (i.e., raw climatology) or a sample of statistically smoothed climatology from historical observations for forecast periods up to 1-year (i.e., resampled climatology). An important feature of the system is that it can also be used to generate hindcast ensembles for system verification and validation. The science underlying the MEFP is described in the *MEFP User's Manual*.

3.1.2 MEFPPE

The MEFP requires parameters that must be estimated for each catchment for which MEFP must generate a forecast ensemble. The parameters are estimated via the MEFPPE, a FEWS explorer plug-in that seamlessly integrates into the CHPS interface and guides the user through the steps of parameter estimation. The general steps to estimate parameters for a catchment are as follows (all actions are performed on a *parameter estimation standalone* which includes MEFPPE):

Steps	Assumptions	Things to watch for
1. Import historical data into CHPS Import historical 6-hour MAP/MAT into the CHPS localDataStore through a datacard import workflow or other mechanism. Convert both time series to a GMT time zone and then convert 6-hour MAT to 24-hour TMIN/TMAX.	6-hour MAP and MAT historical time series must be available. The time series should be the same MAP and MAT time series that are used to drive the standard ESP forecasts at an RFC.	Use the CHPS **Database Viewer** to confirm import and quality control data.

Steps	Assumptions	Things to watch for
2. Make historical data available to MEFPPE Acquire historical MAP/TMIN/TMAX data from the CHPS database and export historical data files in PI-XML for the MEFPPE to use. The files can be exported via a panel within the MEFPPE or can be exported manually by the user prior to running.		The locations available for parameter estimation in the MEFPPE are based on the historical data made available via the exported PI-XML files. The historical time series can be viewed via the MEFPPE.
3. Process historical data Create faster-access binary files containing historical MAP/TMIN/TMAX data, to be stored with the estimated parameters for access during operational ensemble generation.		The processed historical data can be viewed via the MEFPPE. For both precipitation and temperature, the data should be identical to the provided historical data.
4. Create HPC/RFC archive (if needed) The archive of past RFC QPF/QTF (TMIN and TMAX) and corresponding observed values is provided in ASCII text files. It can be created using MEFPPE if the data is in the "vfypairs" table of the archive database. Alternatively, the necessary files can be constructed manually and then imported into MEFPPE.	Archives of past QPF/QTF along with corresponding observed values of several years are available.	The archived QPF/QTF should have been created in the past using the same process that is used for current operational forecasts. The time series can be viewed via the MEFPPE. These archives are necessary to estimate the MEFP parameters for the RFC forecast data source.
5. Acquire GFS/GEFS/CFSv2 reforecast data (one or more, depending on which are used as an input source) Acquire the reforecast data files for the GFS, GEFS, and/or CFSv2 forecast sources as needed via the MEFPPE interface.	MEFPPE acquires the data as needed via SFTP to a site behind the AWIPS firewall.	The time series can be viewed via the MEFPPE. These reforecasts are necessary to estimate MEFP parameters for the GFS, GEFS, and CFSv2 forecast sources.
6. Estimate parameters Specify estimation options and estimate the parameters of the MEFP for whichever forecast sources will be used to generate the ensembles operationally.		Only basic estimation options should be modified by users, initially. Those options are described in the *MEFP User's Manual*. Default values for advanced options should be used in most cases until experience and understanding has been gained with the science of the MEFP.

Steps	Assumptions	Things to watch for
7. Review the parameters and re-estimate if needed Use diagnostic tools in MEFPPE to quality control and review the parameters. Re-estimate parameters with different options if needed.		Make sure the parameters are not unreasonable. For example, the average observed temperature value for any day's parameters should not be unreasonable hot or cold. Make sure the parameters are acceptable before continuing.
8. Accept parameters Using MEFPPE, accept the parameters. This will copy the parameter files from the working area of the MEFPPE to a central directory for access by MEFP executing operationally or while hindcasting.		MEFP parameters are stored as XML and binary files within an gzipped tar (.tgz) file. The parameters are stored by catchment (or location) and by data type (precipitation or temperature).

A more detailed step-by-step procedure for parameter estimation is provided in Section 3.2.3 of the *MEFP User's Manual*. Instructions on how to perform each step using the MEFPPE is provided in the rest of Section 3 of the *MEFP User's Manual*.

An introductory process making use of a **Run All Button** for estimating parameters with minimal interaction is described in the *MEFPPE Configuration Guide* in the context of confirming a successful configuration of MEFPPE.

3.1.3 MEFP Data Ingest Components

Figure 2 provides a diagram of the FEWS transformation modules and OHD model adapters that are part of the MEFP data ingest and forecasting processes. The data ingest components are shown in the upper half and include modules to perform the following actions (note that Step 0 must precede all other steps, but all other steps are performed independently of each other):

Steps	Assumptions	Things to watch for
0. Acquire forcing forecast grids Acquire operational forecasts of GEFS and CFSv2, as needed. The operational forecasts are specified on a gird covering the world and stored in GRIB format files. Remove forecast files older than 30 days.	The FTP process used to acquire these forecasts has been setup.	An FTP script has been provided with the release along with instructions on how to use the script and setup an automated FTP process. See the *MEFP Configuration Guide: Data Ingest Components*. The grids are stored in archived in tar files (.tgz) with subdirectories under the directory specified in the configuration guide: <*ftp_dir*>. The length of the archive is 30 days.

Steps	Assumptions	Things to watch for
Import the GEFS forecast grids Import the GEFS ensemble mean grid files and store them on a subset of the world grid that covers CONUS.	The import workflows for GEFS have been scheduled to run on a daily basis.	Import failure will typically not be detected until grids for that source are used to generate an operational MEFP forecast. When that happens, you may need to import the grid files manually to fill in missing forecasts. The 30-day archive mentioned in Step 0 can be used.
Import the CFSv2 forecast grids and construct location-specific time series XML files Import the CFSv2 grid files and store them on a subset of the world grid that covers CONUS. Interpolate the CFSv2 grids for specific MEFP forecast catchments, yielding a forecast time series for each catchment. Store the time series in a PI-timeseries XML file.	The import workflows for CFSv2 have been setup to run on a 6-hourly basis.	Import, interpolation, or time series export failure will typically not be detected until grids for that source are used to generate an operational MEFP forecast. When that happens, you may need to import the grid files manually to fill in missing forecasts and generate the missing files. The 30-day archive mentioned in Step 0 can be used.

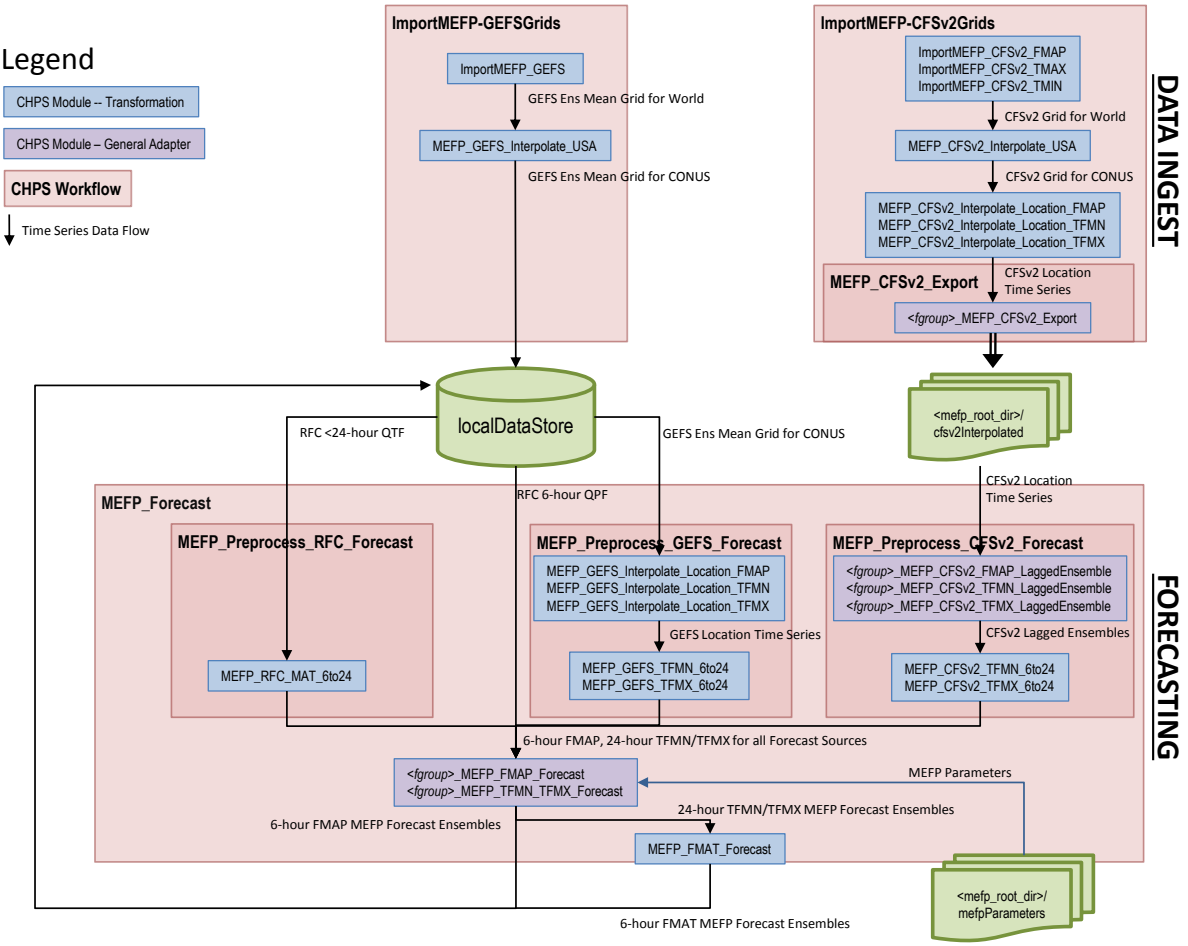

Figure 2: Diagram of the CHPS modules that are part of the MEFP data ingest and forecasting components.

All of the above is accomplished through standard FEWS and CHPS tools and the execution of one new adapter developed for HEFS (but written to be of general use): the TimeSeriesExporterModelAdapter, which is described in Section 4 of the *MEFP User's Manual*. Most of the configuration files required for executing the MEFP data ingest components are provided in the release and any required modifications to those files or other RFC-specific files are described in the *MEFP Configuration Guide: Data Ingest Component*.

Instructions for configuring the MEFP data ingest components, setting up an FTP process to acquire GFS, GEFS, andCFSv2 operational forecasts, and scheduling workflows to import and process the forecasts are provided in the *MEFP Configuration Guide: Data Ingest Components*.

3.1.4 MEFP Forecast Components

The bottom half of the diagram in Figure 2 displays all of the modules that comprise the MEFP forecasting components. The modules perform the following process:

Steps	Assumptions	Things to watch for
1. Prepare RFC QTF data for use in MEFP (if needed) Convert 6-hour QTF FMAT time series to 24-hour minimum (TFMN) and maximum (TFMX) temperature time series.	RFC QTF forecast source is to be used and an appropriate time series is in the CHPS localDataStore.	If RFC QTF is not to be used as a forecast source for MEFP when generating temperature ensembles, then the MEFP_Preprocess_RFC_Forecast workflow can be commented out from the MEFP_Forecast workflow. This is the default for the initial release.
2. Prepare the GEFS data for use in MEFP. Spatially interpolate operational GEFS ensemble mean forecast to compute forecast time series for each catchment and convert the 6-hour TFMN and TFMX forecasts to 24-hour TFMN and TFMX.	See assumptions for the step above, except that the assumptions apply to the GEFS forecast grids and source.	See the things to watch for in the step above, except that the forecast source is GEFS and the workflow that prepares the GEFS data is MEFP_Preprocess_GEFS_Forecast and it is <u>not</u> commented out in the initial release. Also, the GEFS forecast grids specify the ensemble mean; no mean computation is performed.

Steps	Assumptions	Things to watch for
3. Prepare the CFSv2 data for use in MEFP. Construct lagged ensembles (see Section 4.3.7 of the *MEFP User's Manual*) for FMAP, TFMN, and TFMX. Convert the 6-hour TFMN and TFMX lagged ensembles to 24-hour.	See the assumptions for the step above, except that the assumptions apply to the CFSv2 forecast grids. Additionally, the CFSv2 location-specific times series files were successfully interpolated and exported to PI-timeseries XML files.	If CFSv2 forecast source is not to be used as a forecast source for MEFP when generating precipitation or temperature ensembles, then the MEFP_Preprocess_CFSv2_Forecast workflow can be commented out from the MEFP_Forecast workflow. If the CFSv2 forecast grids are not successfully imported or the time series files exported, then errors will occur when executing the preprocessing workflow mentioned above, specifically when generating lagged ensembles. Errors will also occur when executing MEFP if the source is used.
4. Execute MEFP Execute the MEFPEnsembleGeneratorModelAdapter to generate MEFP forecast ensembles.	Inputs are available and exported for use by the MEFP adapter via an exportTimeSeriesActivity for all sources that are to be used in generating forecast ensembles.	Time series that are exported for sources that are <u>not</u> used to be should be removed from the exportTimeSeriesActivity or commented out. Control options for running MEFP are specified in the exportRunFileActivity section of the configuration file. The length of the forecast ensemble is the longest number of forecast days used for any source specified in the control options. It will not be affected by the forecast length associated with the workflow.
5. Create 6-hour FMAT ensembles. Convert the forecast ensembles of 24-hour TFMN and TFMX to a single ensemble of 6-hour FMAT time series.		If the MEFP fails to generate ensembles of TFMN or TFMX, then this step will yield an error in CHPS.

The work in the steps above is accomplished through the use of FEWS transformations and calls to two adapters written specifically for use with MEFP:

- CFSv2LaggedEnsembleModelAdapter: Constructs lagged ensembles from CFSv2 location-specific time series.
- MEFPEnsembleGeneratorModelAdapter: Executes the MEFP algorithm to generate forecast ensembles of FMAP andTFMN/TFMX.

Detailed descriptions of the adapters, including inputs, properties, and outputs, are provided in Section 4 of the *MEFP User's Manual*.

Instructions for configuring the MEFP forecasting components, incorporating them in an ensemble forecast workflow, and generating MEFP hindcasts are provided in the *MEFP Configuration Guide: Forecast Components*.

3.2 *EnsPost*

3.2.1 Description

HEFS EnsPost is a statistical model that post-processes an ensemble of streamflow time series, generating an ensemble that accounts for bias and uncertainty in the hydrologic model, including structure, parameters, initial conditions, and so on. It does not account for input uncertainty, or uncertainty due to meteorological inputs to the hydrologic model, including precipitation and temperature. Thus, it must be applied to a streamflow ensemble forecast that already accounts for input uncertainty. In HEFS, streamflow ensemble forecasts generated using outputs from the MEFP as inputs account for the input uncertainty.

3.2.2 EnsPostPE

HEFS EnsPost depends upon parameters that must be estimated for each segment for which EnsPost will be applied. The parameters are estimated via the EnsPostPE: a FEWS explorer plug-in that seamless integrates into the CHPS interface and guides the users through the steps of parameter estimation. The general steps to estimate parameters for a segment are as follows (all actions are performed on a *parameter estimation standalone* which includes EnsPostPE):

Steps	Assumptions	Things to watch for
1. Identify the module output to which EnsPost will be applied For downstream locations, the module output should comprise the local and routed contributions; i.e., the total flow.	The module must yield streamflow time series at a time step that does not exceed 24-hours and divides 24-hours evenly. The type of data (i.e., CHPS parameterId) is typically either SQIN or QINE.	Observed, historical time series corresponding to module output must be available at a time step no greater than 24-hours in order for parameters to be estimated. See Step 3.
2. Generate a historical simulation for the module to which EnsPost will be applied The resulting historical simulation must be exported to a PI-timeseries file for use in EnsPostPE or stored in the CHPS database and made accessible via the FEWS PI-service. The type of data (parameterId) should be either SQIN or QINE.	Time step of data cannot exceed 24-hours and must divide 24-hours evenly (6 hours or 1 hour is typical).	The historical simulations should originate from the same configuration of CHPS that is used operationally. The historical simulation period should correspond to the period of the record for which EnsPost calibration is desired. The locationId used in exported PI-timeseries file must match that in for the QME historical time series (see below).

Steps	Assumptions	Things to watch for
3. Import historical observed data into CHPS The observed streamflow data should be of type (parameterId) QME (mean discharge). A longer period of record is preferred (e.g. 20+ years). When hindcasting, this should be coordinated with the hindcast period.	Time step of data cannot exceed 24-hours and must divide 24-hours events (6 hours or 1 hour is typical).	It is critical that the time system in the piXML files is correct. Typically, the QME is stored in Data Card format (in local time) and imported into CHPS. When importing the files, the data must be shifted correctly and/or the time system correctly identified in the piXML. The observed stream flow period should correspond to the period of the record for which EnsPost calibration is desired. The locationId used to import the QME historical time series must match that in for the historical simulation (see above). The "time step unit" and "flow units" should be checked for correctness..
4. Evaluate flows Develop annual hydrographs from the QME, and identify dominant seasonality such as wet- or dry- seasons. Calculate the error (simulated flow – observed flow) and plot for the different months/seasons. Calculate basic verification metrics between simulated flows and associated observed flows to inform calibration choices.	Initially, default values may be chosen for the parameters in the EnsPostPE. Over time, experience should lead to more informed choices. Seasons must be formed of consecutive months and the number of months in each season approximately equal (for sample size reasons) unless a strong seasonal error pattern with unequal months in each season exists. Verification metrics and errors should be calculated on the daily time scale (i.e. using QME).	Fewer months per season results in fewer samples in that season. This may increase sampling uncertainty/noise. Configuration files are delivered with the release to allow for using CHPS tools to perform this analysis. This is described in the *EnsPostPE Configuration Guide*.
5. Make historical simulated and observed data available to EnsPostPE Acquire historical QME data from the CHPS database via the FEWS PI-service and export historical data files in PI-XML for the EnsPostPE to use. If not already exported to appropriate PI-timeseries XML file, also acquire the historical simulated streamflow (SQIN/QINE) data and export. The files can be exported via a panel within the EnsPostPE or can be exported manually by the user.		The locations available for parameter estimation in the EnsPostPE are based on the historical data made available via the exported PI-XML files. Use EnsPostPE to view the QME and SQIN data for quality control. Also, the data in the display should correspond to the period of the record for which the parameter estimation of EnsPost is desired

Steps	Assumptions	Things to watch for
6. Set parameter estimation options This includes specifying the seasons and advanced parameter options in the EnsPostPE.	While some parameters, such as choice of seasons, are relatively intuitive, other parameters require an understanding of the technical details and calculation of verification metrics of the EnsPostPE ensembles generated for different options. Without that understanding, the default options are recommended.	The parameters are pertinent to the skill in the simulated streamflow and the basin hydrology. Parameter choices can, therefore, vary with basin.
7. Estimate parameters This estimates the parameter values for the EnsPost.		Observed and simulated streamflows must be available for each location.
8. Review the parameters Once the parameter values have been estimated, they should be checked for plausibility and acceptance.	The *EnsPostPE User's Manual* provides some guidance, but diagnostic information will be improved with future DRs.	Visualize the parameters using the EnsPostPE GUI. Typically the regression coefficient, b, is high for high flow values. Parameters for all months within a single season should have the same values.
9. Verify EnsPostPE output This is an optional step, but recommended. The EnsPostPE provides corrected ensembles that can be verified directly. By comparing these ensembles for different options of the PE, both against the observed flows and the simulated flows, the options used in the PE may be better informed	This step is time consuming and requires knowledge of how to evaluate the performance of the EnsPostPE ensembles (e.g. using the EVS). However, it might be considered for critical locations or where time permits	
10. Accept parameters Using EnsPostPE, accept the parameters. This will copy the parameter files from the working area of the EnsPostPE to a central directory for access by EnsPost executing operationally or while hindcasting.		EnsPost parameters are stored as ASCII text files within a gzipped tar (.tgz) file. The parameters are stored by location and by historical simulation data type (SQIN or QINE).

A more detailed step-by-step procedure for parameter estimation is provided in Section 3.2.3 of the *EnsPost User's Manual*. Instructions on how to perform each step using the EnsPostPE is provided in the rest of Section 3 of the *EnsPost User's Manual*.

An introductory process making use of a **Run All Button** for estimating parameters with minimal interaction is described in the *EnsPostPE Configuration Guide* in the context of confirming a successful configuration of EnsPostPE.

3.2.3 EnsPost Forecast Components

There is only one forecast component used for HEFS EnsPost: the HEFSEnsPostModelAdapter. Module configuration files must be created to call the adapter for each segment to which EnsPost will be applied, providing observed streamflow and an ensemble to post-process to the adapter, which in turn generates a post-processed ensemble. Id-mapping and run file properties can be used to specify the parameter file EnsPost to use. The EnsPost modules must then be called from the existing workflows that generate streamflow ensemble forecasts to be post-processed for those segments.

3.3 *EVS*

3.3.1 Description

The Ensemble Verification System (EVS) is designed to verify ensemble forecasts of continuous numeric variables, such as temperature, precipitation, streamflow and river stage. The EVS can be applied to forecasts from any number of geographic locations (points or areas) and issued with any frequency and forecast lead time. It can also aggregate forecasts in time, such as daily precipitation totals based on hourly forecasts, and can aggregate verification statistics across several discrete locations.

A verification study with the EVS is separated into three stages, namely: 1) Verification; 2) Aggregation; and 3) Output. In the Verification stage, one or more Verification Units (VUs) are defined. Each VU comprises a set of forecasts and verifying observations for one environmental variable at one geographic location (e.g. one river catchment). The ensemble forecasts and observations are provided in a PI-XML or ASCII format. Currently, the EVS cannot read directly from the CHPS database. The Verification stage also requires one or more verification metrics to be selected. The forecasts and observations are then paired by forecast lead time and the verification metrics computed. The results are written to the Output dialog, where the metrics can be plotted in an internal viewer or written to file in a variety of graphical formats or in XML. The Aggregation stage allows for the averaging of verification statistics across multiple VUs. Once an EVS project has been defined, that project may be executed independently from the CHPS, either from the command line or within the EVS GUI. It can also be executed within the CHPS, at the end of a CHPS hindcast workflow, using the EVS-to-CHPS model adapter.

The verification metrics in the EVS comprise both deterministic metrics, which verify the ensemble average forecast, and probabilistic metrics, which verify the forecast probabilities. A full description of the EVS, including the verification metrics, is available in the *EVS User's Manual*.

The general steps to conduct a verification study with the EVS are summarized below. They assume that an archive of forecasts and corresponding observations is available. This may be produced by hindcasting (see Section 3.4) or by archiving operational forecasts and observations, assuming the operational models and configurations are relatively stable over the archive period.

Steps	Assumptions	Things to watch for
1. Collect data Collect data for each variable (e.g. precipitation), forecast location and scenario (e.g. streamflow with EnsPost) to be verified. Are "sufficient" data available? Generally hindcasting will be used to provide a long record (5+ years), ideally 20+ years.	Data are available. Precise data needed depends on what is being verified. At least forecasts and observations for pairing. Other variables may be used to "condition on" (e.g. investigate quality of precipitation forecasts when temperature is below freezing). If evaluating skill, the forecasts for the baseline are also needed (e.g. "ESP").	If forecasts and observations are measured at different times or cover different control volumes (e.g. 6-hours versus daily), a strategy is needed to pair the forecasts and observations
2. Create EVS project file Use a template or existing EVS project file or start from scratch (create a new project).	When running the EVS in standalone mode, no CHPS configuration is required. When running in CHPS mode, the project file is zipped and placed in an appropriate location for CHPS to access (see Table 5).	Generally, it is simpler to start with an existing EVS project file or "template" because many verification studies are similar (e.g. just applied to different locations). The HEFS Development Team can provide templates.
3. Add Verification Units A Verification Unit (VU) is required for each variable, location and scenario to verify. The VU identifies the location of the data, the time-scales to be verified, and the metrics to compute, among other things. An EVS project file generally contains several VUs.		Verification with the EVS can be very time-consuming, depending on the size of the dataset being verified. Thus, careful thought about the aims of the verification study can save considerable time overall.
4. Subset data if needed For each VU, it is possible to apply conditions to extract subsets from the overall dataset. For example, the data may be broken into particular seasons and verification conducted separately for each season.	That sufficient data are available to compute the verification metrics for the subsets of data identified, otherwise the sampling uncertainties may be large.	Large sampling uncertainties (i.e. "noisy" or misleading verification results) when using small sample sizes.
5. Configure metrics In general, it is necessary to compute several verification metrics. One metric cannot provide a complete picture of forecast quality. However, some metrics are more suitable to particular problems. The EVS metrics are grouped into single-valued, ensemble and skill metrics.	Requires some preparatory thought about the verification metrics that should be computed (i.e. what is the aim of the verification study?) and any thresholds that should be used to compute them (e.g. flood thresholds). Is it necessary to quantify the sampling uncertainty (e.g. small sample size)?	Consider the computational time when choosing a large number of metrics and computing them at a large number of thresholds.
6. Run the verification Two steps are conducted by the EVS when running a VU: 1) the forecasts and observations are paired; and 2) the verification metrics are computed with the paired data.	That forecasts and observations are available at the same times and for the same accumulation volumes, otherwise the pairs will not be computed. That "sufficient" data are available for verification (generally several years).	Verification may be time consuming and CPU/memory intensive. When applying the EVS to large datasets, it may be necessary to increase the maximum memory allocated to the EVS **before** start-up (see start-up options in the *EVS User's Manual*).

7. Check verification pairs All verification results from the EVS reflect the verification pairs that were computed by the EVS. It is, therefore, critical to check some of these pairs against the raw data before relying on the results from the EVS.	That the pairing has been conducted correctly. Without correct pairing of the forecasts and observations, the verification results will be meaningless.	The verification pairs are stored in an XML format with times in UTC. The pairs reflect any aggregation requested when defining the VU needed to do the pairing (e.g. aggregation of 6-hourly forecasts to pair with 24-hourly observations).
8. Display/interpret results The verification results may be viewed outside of the EVS (using the written outputs) or with the interactive viewer inside the EVS GUI. Interpreting the verification results requires time and practice. The EVS user's manual provides some guidance on the meaning of the different metrics. Some metrics are more intuitive than others. For example "skill scores" show the relative quality of one forecasting system (e.g. MEFP) given a baseline (e.g. climatology).	Some familiarity with the verification metrics available in the EVS and with ensemble verification more generally. An awareness of the application and audience for the verification results.	Some metrics provide relative measures of quality and others are expressed in forecast units (e.g. CFS for streamflow). With the latter, some care is needed, because the "meaning" of these units will vary substantially between locations. Take care with interpreting results that are based on small sample sizes. The numerical (XML) outputs from the EVS provide sample sizes. Sampling uncertainties can also be evaluated explicitly, but this is time consuming.

3.3.2 EVS-to-CHPS model adapter

In general, it is simpler to run the EVS outside of the CHPS, as a stand-alone application, because the EVS requires extensive configuration to produce the project file used to conduct a verification study. Once the configuration is done, however, execution is straightforward and does not rely on the CHPS, because the EVS does not read from the CHPS database. However, in some cases, it may be desirable to run the EVS inside of the CHPS after a hindcast run, i.e. at the end of a hindcast workflow where the hindcasts have been exported to files. For that purpose, a model adapter has been made available that allows for verification statistics to be computed as part of a CHPS workflow:

- EVSToFEWSAdapter: Compute verification statistics based on a previously saved EVS project using data exported during a hindcasting workflow execution in CHPS.

3.3.3 Hindcasting

Critical to performing a verification study is the ability to generate hindcasts using components of the HEFS, including MEFP, hydrologic models, and EnsPost, as it is those hindcasts that will be verified. Generating hindcasts can be done using the batch-run option of the CHPS **Manual Forecast Dialog**. It requires constructing an end-to-end (MEFP through EnsPost) workflow, if one does not already exist, which will then be the one that is executed for hindcasting. It also requires some preparation work for MEFP and adding time series exports to the hindcasting workflow. For detailed information, see the *Streamflow Hindcasting Cookbook*.

The following general steps must be performed when generating a hindcast:

Steps	Assumptions	Things to watch for
1. Collect historical data Collect historical data (including MAP, MAT, MAPE, and QME) for hindcast forecast location. MAP, MAT, and MAPE data are required to generate warm states for the hindcasting run. QME data is required to run EnsPost.	Data are available.	Generally a long record (5+ years), ideally 20+ years, of data is required. When importing historical MAP, MAT, and QME data, an offset of -6 hours should be added to make sure the data ingested will be in the right timezone (GMT-6). If datacard files are used as a source for the historical data, then it must follow specific rules outlined in the *MEFPPE Configuration Guide*.
2. Prepare MEFP for hindcasting Add a hindcasting flag to the configuration indicating that MEFP will run in hindcast mode. Turn off all pre-processing workflows associated with forecast sources.	The parameters for MEFP have been appropriate estimated, meaning that reforecasts and archived forecasts needed for hindcasting are stored in the parameter file.	The hindcasting flag must be set in every module that executes the MEFPEnsembleGeneratorModelAdapter. This process can be simplified if using a global property. See the *MEFP Configuration Guide: Forecast Components*.
3. Prepare EnsPost for hindcasting If EnsPost has already been configured for the operational workflow, then only the time series (the raw streamflow hindcasts to be post-processed) input to EnsPost need to be updated in the EnsPost module instance. Particularly, the ensemble ID of the new time series should be specified.	The parameters for EnsPostPE have been appropriate estimated and stored in the parameter file.	
4. Configure an end-to-end parent workflow If one does not already exist, create a parent workflow that executes the entire streamflow process end-to-end. This must include the MEFP_Forecast workflow and the EnsPost workflow. If one wants to export MEFP hindcasts, streamflow hindcasts, and EnsPost-processed hindcasts, corresponding export workflow and module instances should be created and add to the end-to-end workflow.	If MEFP is used, then it has been properly configured to execute. Further, the ensemble streamflow forecast workflow has been configured to use the MEFP output as input to the streamflow forecast process. If EnsPost is used, then it has been configured and added to the parent workflow.	The forecast workflow for the upstream basins must be put before which of the downstream basins.

Steps	Assumptions	Things to watch for
5. Configure an EVS workflow if needed Create a module to run an EVS project and workflow to execute it.		The location identifier and variable identifier in the EVS project file should match those in the PI-XML files exported by CHPS. All new workflow and module files should be registered appropriately.
6. Generate warm states Use the **Manual Forecast Dialog** to execute an update states workflow in order to generate states for the dates in the past for which hindcasts are to be generated.		Warm up the model (using at least a 2-year period) before running the UpdateStates workflow. Need to make sure the warm state search period ends on day 0 for the end-to-end parent workflow. This can be defined in *WorkflowDescriptors.xml*. If MEFP is used to drive the hindcasts, then the dates for which hindcasts can be generated is dictated by the availability of reforecasts for the used forecast sources. See the *MEFP Configuration Guide: Forecast Components* for hindcast date restrictions.
7. Generate streamflow hindcasts Use the **Manual Forecast Dialog** batch run option to generate hindcasts for the dates in the past for which warm states were generated.	The warm were successfully generated for each hindcast date.	Define the forecast length appropriately (not longer than the length of the MEFP output). If the MEFP is used as a source of input for the streamflow forecasts and Step 1 was not performed above, then errors will occur. Errors in preprocessing workflows can be ignored, but errors in MEFP because the hindcasting flag was not set cannot be ignored. Be sure to configure EnsPost as part of the streamflow forecast workflow if EnsPost is to be used.
8. Confirm hindasts were created. Use the database viewer to view created hindcasts and confirm that the time series to be verified later were created.		
9. Run EVS workflow If running an EVS workflow, execute a verification run to produce verification products for streamflow hindcasts at test locations		Define the parameter ID appropriately (corresponding to the variables to be verified) in EVS module files.

4 What Next?

4.1 *HEFS System*

[Most of this subsection is copied from the HEFS ConOps.]

The HEFS will become a supported part of the CHPS baseline in early 2015. Until then, i.e. for the HEFS rollout, there are a few system options for RFCs to run the HEFS. When selecting an option, the RFC should weigh whether they want to isolate HEFS from their operational CHPS during their learning HEFS implementation versus the added effort of maintaining a separate set of configurations. RFCs planning to use their CHPS Dev/Test system for the rollout should ensure that system is working.

The recommended option is for the HEFS to be configured on both the CHPS operational and CHPS Dev/Test systems and run the HEFS executables only on the CHPS Dev/Test system. The systems are synched, and the operational CHPS is isolated by restrictive file/directory permissions on the operational CHPS. For instructions to restrict permissions, see Fogbugz 1371:

http://schuylkill.nws.noaa.gov:7069/default.asp?1371

Once configuration and testing is complete, the permission restrictions may be removed if desired.

Another option is to configure and run HEFS only on the CHPS Dev/Test system with a live feed from the operational CHPS. In this scenario, the CHPS Dev/Test (running HEFS) is not synched with the operational CHPS. The advantage is that the HEFS is kept completely separate from the operational CHPS. However, this will temporarily necessitate that CHPS configuration changes be made to both the operational CHPS and CHPS Dev/Test (with HEFS) configurations. Maintaining both configurations will no longer be necessary after HEFS configurations are moved into the operational CHPS configurations (via a merge).

The third option, if the RFC is willing to accept the risk, is that the HEFS can be configured on an RFC's operational CHPS. This option has the advantage of the RFC not having to maintain two (CHPS and HEFS) configurations. This is risky because during the rollout the remaining RFCs will have no or limited experience with HEFS, especially the impact of HEFS on their run-times (of existing non-HEFS workflows) and database storage.

4.2 *Installing HEFS*

After reading this manual, the next step is to install and configure the HEFS components that are to be used. Begin by identifying which components you want to use as part of operational forecasting. Then, identify a first segment for which to install the desired HEFS component(s) and, if MEFP is to be used, identify all catchments employed by that segment; the segment is referred to as the *installation segment* and catchments as *installation catchments*.

If the MEFP is to be used, then do the following:

1. Configure the MEFP data ingest components, setting up the FTP process and scheduling the data ingest workflows. See the *MEFP Configuration Guide: Data Ingest Components*. In order to run MEFP with any of the gridded forecast sources, GFS, GEFS, or CFSv2, the data ingest process must be working and data must be in the CHPS database (for GFS and GEFS) or on the file system (for CFSv2).
2. Configure and execute the MEFPPE to acquire parameters for the *installation catchments*. The parameter estimation process is done in a *parameter estimation standalone* that can be reused for EnsPostPE. See the *MEFPPE Configuration Guide*, which includes instructions for a quick process to estimate parameters (to get started; full step-by-step instructions are in Section 3 of the *MEFP User's Manual*).
3. Configure the MEFP forecast components. Confirm that the MEFP generates reasonable output time series. See the *MEFP Configuration Guide: Forecast Components*.
4. Add the MEFP forecast components to an ensemble streamflow forecast workflow and confirm that it executes. See the *MEFP Configuration Guide: Forecast Components*.
5. (Optional) Create an end-to-end (MEFP through streamflow) workflow and add it as a scheduled workflow to generate MEFP-based ensemble forecasts on a daily basis.
6. Configure the HEFS Graphics Generator products following instructions in the *HEFS Graphics Generator Products Installation Guide*.
7. Tweak the configurations and add additional forecast groups, segments, and catchments as desired.

If the EnsPost is to be used, then do the following:

1. Identify an ensemble forecast workflow to which EnsPost will be added. This could be the MEFP-based streamflow ensemble workflow setup for the MEFP, above.
2. Configure and execute the EnsPostPE to acquire parameters for the *installation segment*. The parameter estimation process is done in a parameter estimation standalone that can be reused for MEFPPE. See the *EnsPostPE Configuraiton Guide*, which includes instructions for a quick process to estimate parameters (to get started).
3. Configure the EnsPost forecast components and add it to an appropriate ensemble streamflow workflow, such as that created for the MEFP, above. Confirm that EnsPost generates reasonable time series. See the *EnsPost Configuration Guide*.
4. Configure the HEFS Graphics Generator products following instructions in the *HEFS Graphics Generator Products Installation Guide*.

If the Graphics Generator is to be used to generate HEFS products, then do the following:

1. Configure the HEFS Graphics Generator products. See the *HEFS Graphics Generator Product Installation Guide*.
2. If the products are to be created on a routine basis, create modules to generate the products and a workflow to call those modules. Schedule the workflow so that it executes after all required input time series are generated, or add it to the end of that workflow. See the *HEFS Graphics Generator Product Installation Guide*.

Appendix A: Limitations

The HEFS is an operational system and is subject to regular enhancements. These include phased enhancements and bug fixes, which are based on scientific evaluation and software testing. The phased enhancements are implemented in "Development Releases" (DRs). Scientific evaluation requires hindcasting and verification, which are time-consuming and resource intensive. Also, the research-to-operations transition of the HEFS will lead to several novel applications that may require further testing and evaluation. The HEFS Version 1.0 has several known limitations, of which some will be addressed in the planned DRs. The main limitations are summarized below.

Limitation	Potential impact	Plans to address
Limited functionality for quality controlling the HEFS components, including the calibration of the MEFP and EnsPost (i.e. the PEs) and real-time application, i.e. how to identify problematic forecasts	Difficult to tune parameters of the MEFP and EnsPost for particular applications. Relies more on hindcasting and verification, which is time consuming. In real-time application, "problem" forecasts may be difficult to screen	Tools have been added for the quality control of historical data and reforecasts/archived forecasts used during parameter estimation. Additional tools are required to guide users through the process of selecting parameter estimation options and judging the quality of estimated parameters. Real-time screening of HEFS forecasts is a potential future enhancement after HEFSv1
Inability to **explicitly** account for some sources of uncertainty, notably in hydrologic model states and parameters, and in observations. Instead, their effects are accounted for indirectly (by EnsPost)	More difficult to isolate particular problems or sources of uncertainty to be addressed by enhancements. Lumping of different sources of uncertainty runs risk that some aren't properly accounted for and relies heavily on (quality of) observed flow and EnsPost. Reliance on manual MODs rather than data assimilation (DA) could lead to inaccurate model states and improper accounting for uncertainty	No plans to address inaccurate model states until at least HEFSv2 through DA.
Limitations of the simple assumptions made by the HEFS components, notably MEFP and EnsPost, when addressing complex hydrometeorological / hydrologic conditions	Many specific instances, but key examples include river regulations, extreme events and cases where the residuals of the fitted models (MEFP and EnsPost) are not normal. Also, the space-time modeling adopted by the MEFP and EnsPost is quite simplistic	This will be addressed in guidance for the specific components.
Limited sources of raw forcing forecasts. Currently limited to RFC/HPC, GEFS, and CFSv2	Failure to accommodate potentially valuable forcing information, such as forecasts from the SREF. However, this ideally requires a suitable archive of hindcasts	Additional sources of forcing information will be addressed after HEFSv1, as needed.

Limitation	Potential impact	Plans to address
Limited flexibility of the science algorithms. Inability to choose an algorithm for a particular situation, based on guidance	Some scope for "tuning" the MEFP and EnsPost, but limited scope for changing the underlying modeling approach to suit the application. When the assumptions of the MEFP and the EnsPost are not fully met, there are no alternatives to apply	No plans to increase the flexibility of the science algorithms. The provision of a tool box of techniques for the forcing and streamflow will be explored in future
Limited pre-defined products or templates for communicating the outputs from the HEFS. While the GraphGen and the EVS are both flexible, templates are also needed for HEFS products	Potential confusion about how best to communicate the outputs from the HEFS or lack of consistency between RFCs (some of which may be justified)	This is an ongoing effort and will be improved by knowledge of how the HEFS is being applied in practice. This is not, primarily, a software issue, but related to the development of templates and guidance for applying the GraphGen (operationally) and the EVS (for hindcasting)
Limitations of the underlying hydrometeorological and hydrologic models used in the CHPS	This is broad problem. Examples include limitations of the lag/K routing approach, inability of the raw forcing models to capture convection, difficulties in calibrating Snow-17 etc.	No specific plans to address these limitations. In terms of routing, the three-parameter Muskingum model has been investigated and may be considered in future
Limited hindcasting and verification of the HEFS components, as well as "end-to-end" applications	Limited insight into the quality and skill of the HEFS ensembles under varied conditions, including situations where the HEFS performs less well. Limited guidance on how to apply the HEFS in practice	This is currently being addressed through three phases of hindcasting and verification, mainly focused on the different sources of raw forcing information (via MEFP) and the application of EnsPost. The hindcasting and verification will be used to develop improved guidance and build trust in the HEFS. However, this only constitutes a preliminary effort, focused on a limited range of basins and fixed configuration options; further hindcasting and verification will be required in future
Limited ability to plot large datasets in GraphGen. For example, inability to plot hourly data for ~40 ensemble members for more than ~240 days	Reduced scope for visualization of long-range predictions at an hourly timestep	In practice, it should be possible to visualize the long-range forecasts at reduced frequency (e.g. 6-hourly or daily)
MEFP Parameters must be re-estimated whenever a model underlying a forecast source, such as the RFC QPF/QTF, GEFS, or CFSv2, is changed.	Significant time and resources required for OHD to acquire new reforecasts that include the changes to the models and reformat those reforecasts into the format MEFP expects. Significant time required for RFCs to re-estimate the MEFP parameters once the reforecasts are available.	Options will be discussed with NCEP and OHD management.

Limitation	Potential impact	Plans to address
EnsPost Parameters must be re-estimated for a segment whenever a change is made to the hydrologic models (SAC-SMA, SNOW17, LAGK, etc.) for that segment. This includes modifying parameters of a model or adding/removing module to/from the segment.	Significant time and resources needed to generate new historical simulations and re-estimate parameters.	Recommend HEFS to be run in an isolated environment with limited hydrologic model changes.
Discrepancies between offline applications of the HEFS for hindcasting and verification versus online use for operational forecasting. For example, model states may be adjusted in real-time through manual modifications or sources/sinks introduced using real-time information (e.g. on flow diversions).	This has several potential impacts, notably: 1) potential discrepancies between the calibration of the EnsPost versus real-time application (e.g. accounting for certain sources/sinks in one but not the other); and 2) potential discrepancies between the evaluation of the HEFS through hindcasting and verification versus real-time application, resulting in misleading guidance on forecast quality from hindcasting.	A comprehensive Concept of Operations (CONOPS) is currently being developed, in order to guide the calibration, hindcasting/verification and real-time application of the HEFS at the RFCs. The CONOPS will address the potential discrepancies that could arise between offline use (calibration and hindcasting/verification) versus real-time application and make recommendations to mitigate these. In future, there is a need to archive the operational forecasts from the HEFS, as well as conduct hindcasting studies.
Difficulties in calibrating and applying the HEFS to regulated rivers, where some regulations may be unknown, whether historically or in real-time (see above also).	The potential impacts include the improper calibration of the EnsPost for regulated rivers, where historical simulations and observations may differ in the regulations included or differ from the assumptions made in real-time (e.g. about flow diversions). Also, even where flow records are consistent, the EnsPost is not designed for complex time-series involving abrupt changes in statistical behaviors, potentially caused by regulations.	A comprehensive Concept of Operations (CONOPS) is currently being developed, in order to guide the calibration, hindcasting/verification and real-time application of the HEFS at the RFCs. The CONOPS will address the potential issues with calibrating and applying the HEFS in regulated rivers and the need to archive real-time regulations. Further hindcasting and verification is needed to establish the quality of the HEFS forecasts for a range of regulated rivers and forecasting scenarios. Finally, there is a need to leverage and evaluate emerging techniques for handling river regulations in an ensemble forecasting context, such as the “regulated ESP” approach of the Hydrologic Research Center.

Glossary of Terms, Acronyms, and Abbreviations

Aggregation/Disaggregation – forming larger or smaller control volumes, respectively

ASCII – American Standard Code for Information Interchange

Bias – A systematic difference between an estimate of some quantity and its "true" value (generally meaning observed)

BS (Brier Score) – the average squared deviation between the predicted probabilities that a discrete event occurs (such as flooding) and the observed outcome (0 or 1)

Calibration – A process of estimating model parameters based on observations and corresponding (raw) predictions. In hydrologic model calibration, the calibrator is able to modify the parameters directly in order to tweak model performance against a baseline. In pre- and post-processing, calibration has a second meaning, namely to correct probabilistic biases in ensemble forecasts by increasing their reliability

Canonical Event – a partitioning of time scales in order to account for the varying skill of the different forcing inputs to MEFP (e.g., RFC QPF/QTF, GFS, and CFSv2)

CDF (Cumulative Distribution Function; see Probability Distribution) – describes the probability of a variable being less than or equal to a value; i.e. Prob($X \leq x$).

CFS/v2 (Climate Forecast System) – Climate Forecast System. A fully coupled model representing the interaction between the Earth's oceans, land and atmosphere that can generate a forecast for 45 days, a full season, or 9 months. See also: http://cfs.ncep.noaa.gov/

CHPS (pronounced "chips") – Community Hydrologic Prediction System.

Climatology – The science that deals with the phenomena of climates or climatic conditions. Climatology also refers the historical record of observations (e.g. mean areal averages of actual temperature and precipitation) used to drive a model

Correlation Coefficient – Pearson product-moment correlation coefficient. The covariance of two variables divided by the product of their standard deviations. A degree of linear association between two variables, with -1 and 1 denoting perfect negative and positive association, respectively, and 0 denoting the absence of a linear association (but not necessarily a non-linear association)

CPC – Climate Prediction Center)

CRPS (Continuous Ranked Probability Score) – The integral square difference between a forecast probability distribution and the observed outcome. It is typically averaged over many such cases (known as the "mean CRPS")

DA (Data Assimilation) – A procedure for updating model states (and possibly other variables) with recent observations, thereby improving forecasts.

Deltares (formerly Delft) – Netherlands company that developed FEWS which is "wrapped by" CHPS

Disaggregation – (see aggregation/disaggregation)

DR – Development Release

Ensemble Forecast – A collection of equally likely predictions of the future states of a hydrologic system, based on sampling of the different sources of uncertainty and propagating them through a hydrologic modeling system (such as CHPS). An "ensemble trace" comprises two or more forecast lead times

EnsPost (Ensemble Post-Processor) – A software tool and statistical technique that accounts for hydrologic uncertainties and biases separately from the forcing uncertainties and biases

EPP3 (Ensemble Preprocessor) – A (Fortran) pre-cursor to MEFP (Java)

ESP (Ensemble Streamflow Prediction) – In NWS operations, this has the specific meaning of forcing the NWS River Forecast System with a sample of observations from the same dates in the past, i.e. climatological forcing. Some RFCs have augmented the original ESP algorithms to account for additional information

EVS – Ensemble Verification System

FEWS (Flood Early Warning System) – Developed by a company in the Netherlands, Deltares (formerly Delft) and written in Java. See also CHPS

Forcings – The model inputs (e.g., precipitation and temperature) that drive/"force" a hydrologic model

Forecast LeadTtime – The difference between the Forecast Valid Time and the Forecast Issue Time

Forecast Time – see system time.

Forecast Valid Time – The time at which a forecast is valid

Fortran – A general-purpose, procedural, imperative programming language

GEFS (Global Ensemble Forecast System) – an enhanced version of the GFS that produces ensemble forecasts

GFS (Global Forecast System) – One of the operational forecast models run at NCEP. The operational GFS is run four times daily, with forecasts out to 384 hours. The GFS was also "frozen" in 1997 (the "frozen GFS") and used to generate hindcasts (i.e. retrospective forecasts) beginning in 1979, which can be used to calibrate the MEFP. However, operational

forecasting with the GFS is no longer possible, as the frozen model was discontinued from operational use in August 2013.

Grib File – A binary file format designed to store large amounts of gridded data; used for GFS, GEFS, and CFSv2 operational forecasts

GUI – Graphical User Interface

HEFS – Hydrologic Ensemble Forecast Service

Hindcast – a retrospective forecast or reforecast. A forecast where the issue time (T0) is in the past, based upon the conditions at the chosen T0, but using a current model (which may not have been available on the original forecast date). Reforecast is a term frequently used for weather models

HPC – Hydrometeorological Prediction Center

MAP – Mean Areal Precipitation over a basin/watershed

MAT – Mean Areal Temperature over a basin/watershed

MEFP (Meteorological Ensemble Forecast Processor) – A software tool and statistical technique that produces ensemble forecasts of temperature and precipitation using (single-valued) operational forecasts from NWP models. The forecast spread is derived from historical information about forecast errors

NCEP – National Centers for Environmental Prediction

Normal/Gaussian Distribution – A simple, theoretical, probability distribution with two parameters (mean and standard deviation). The multivariate normal distribution, which describes several forecast times, locations or variables, is completely defined by a vector of means (one for each variable) and a covariance matrix

NQT (Normal Quantile Transform) – A transformation made to a data sample so that it follows a normal probability distribution (i.e. so that the histogram of values would appear normal)

NWP – Numerical Weather Prediction

NWSRFS (National Weather Service River Forecast System) – Replaced by CHPS

Parameter Estimation – A process of estimating model parameters based on observations and corresponding (raw) predictions. In parameter estimation, the user can modify control options that affect the estimated parameters, but cannot modify the parameter directly (this distinguished it from calibration; see calibration). Calibration is sometimes used synonymously with parameter estimation, but we avoid the term calibration here as it has several different meanings.

PDF (Probability Density Function; see Probability Distribution) – a continuous function that is used in the calculation of the CDF; it is a continuous version of the Probability Mass Function (see below).

PEDTSEP – A sequence of letters that identifies a type of data; in the Standard Hydrologic Exchange Format (SHEF), different types of data are keyed by a seven-character parameter code represented by the character string "PEDTSEP". This string is broken down as follows:

PE = Physical Element (precipitation, gage height, temperature, etc.)
D = Duration Code (instantaneous, hourly, daily, etc.). The duration code character (D) combined with the physical element (PE), describe the vast majority of observed hydrometeorological data. The duration code describes the period to which an observed or computed increment applies
T = Type Code (observed data, forecast data, etc.)
S = Source Code (further refines the type code which may indicate how data was created or transmitted
E = Extremum Code (maximum value, minimum value, etc.)
P = Probability Code (Chance value is at/below the specified value, e.g., 90%, 10%)

Example: 6-hour precipitation would be encoded PPQ, where PPQ represents incremental precipitation and the PPQ represents a 6-hour duration

For more, see: https://ocwws.weather.gov/intranet/whfs/SHEF/Explain_duration2.shtml

PI – Published Interface

PoP (Probability of precipitation) – The probability that a non-zero precipitation amount will occur.

Probability Distribution – a function that describes the probability of each possible event associated with a random variable, usually quantified by a cumulative probability distribution. A discrete random variable, such as the possibility of flooding, is described with a discrete cumulative probability distribution and a probability mass function. A continuous random variable, such as temperature, is described with a continuous cumulative probability distribution and probability density function. A mixed random variable, such as precipitation, is described with a mixed probability distribution (i.e. precipitation has a discrete component associated with no precipitation and a continuous component associated with positive precipitation).

Probability Mass Function (see Probability Distribution) – for a discrete variable, it describes the probability of the variable taking on a value; i.e., Prob(X = x).

RAX (RFC Archive Database) – An archive of RFC forecasts and observed data stored in a Postgres database

Reforecast – See Hindcast.

RPS (Ranked Probability Score) – An extension of the Brier Score to several discrete probability categories (such as low, medium and high flows). Extension to all possible categories of a continuous variable is equivalent to the CRPS

SHEF (Standard Hydrologic Exchange Format) – A standard ASCII format for exchanging data at the National Weather Service (NWS)

Simulation – A hydrologic prediction based on observed temperature and precipitation (as distinct from a forecast, which comprises forecast inputs).

Skill – The fractional improvement of one forecasting system relative to a baseline. The measure used for skill could vary (e.g. the mean error of one system relative to another)

SOA (Service Oriented Architecture) – An approach to developing software that emphasizes developing software in the form of interoperable services.

SREF (Short-Range Ensemble Forecast) – An NCEP model/system that issues short-range ensemble forecasts

System Time (T0, Forecast Time) – The date/time at which a forecast is initiated.

T0 – see system time.

UTC – Coordinated Universal Time, also known as Zulu time (Z-time) and synonymous with Greenwich Mean Time (GMT). Forecasts from the HEFSv1 are issued daily at 12Z

XEFS **(Experimental Ensemble Forecast System)** – The experimental precursor to HEFS

XML **(eXtensible Markup Language)** – XML is a markup language that defines a set of rules for encoding documents in a format that is both human-readable and machine-readable. See also:

http://www.weather.gov/glossary/
http://www.crh.noaa.gov/dtx/glossary.php
http://www.nws.noaa.gov/mdl/synop/acronyms.php

Also, consider using the search feature on any of the above web sites.