

The Design of a Modern Information Technology Infrastructure to Facilitate Research-to-Operations Transition for NCEP's Modeling Suites

LIGIA BERNARDET, LAURIE CARSON, AND VIJAY TALLAPRAGADA

NOAA/NCEP employs a number of NWP models to provide operational guidance to its service centers and to the NWS field offices. The Global Spectral Model (GSM) is used for the global applications and three other models are used for limited-area applications: the Nonhydrostatic Mesoscale Model (NMM), the Advanced Research WRF (ARW) model, and the Nonhydrostatic Multiscale Model on the B grid (NMMB). Note that while the operational ARW and NMM dynamic cores employ the WRF

framework, NMMB employs the completely distinct NOAA Environmental Modeling System (NEMS) framework, and the GSM is gradually being transitioned to use NEMS. In addition to the various models and frameworks used, NCEP runs a large number of modeling suites, defined here as NWP systems with multiple workflow components (data assimilation, forecast model, postprocessing, etc.), assembled for specific applications (Fig. 1). Each suite depends on several workflow components

besides a myriad of operations related to retrieving and staging input files, creating output directories, archiving output, purging disks, etc., schematically shown in Fig. 2.

Much attention has been given recently to reviewing the requirements for NCEP numerical guidance, with an eye toward designing the most effective and integrated modeling strategy connecting projects across NOAA. For example, a unified Next-Generation Global Prediction System (NGGPS) is being developed to generate an operational forecast guidance

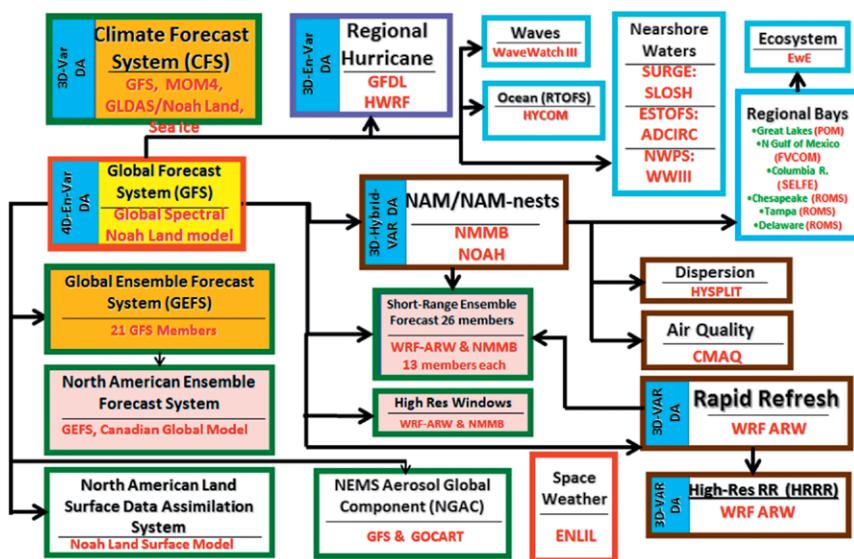


FIG. 1. NWP applications run operationally at NCEP as of March 2016.

AFFILIATIONS: BERNARDET—NOAA OAR/ESRL Global Systems Division, and University of Colorado/Cooperative Institute for Research in the Environmental Sciences, and Developmental Testbed Center, Boulder, Colorado; CARSON—NCAR Research Applications Laboratory, and Developmental Testbed Center, Boulder, Colorado;

TALLAPRAGADA—NOAA NWS/NCEP Environmental Modeling Center, College Park, Maryland
CORRESPONDING AUTHOR: Ligia Bernardet, ligia.bernardet@noaa.gov

DOI:10.1175/BAMS-D-15-00139.1

©2017 American Meteorological Society

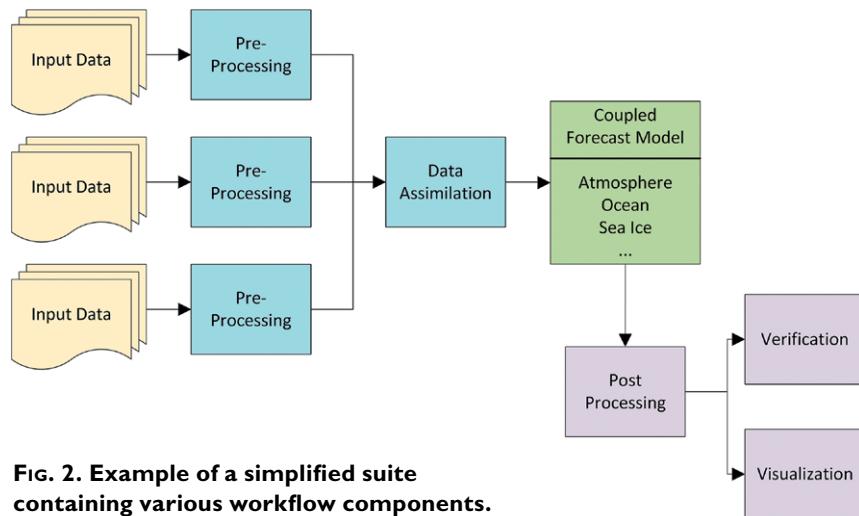


FIG. 2. Example of a simplified suite containing various workflow components.

system for possibly the next few decades. This system will include forecast applications for various scales (from weather prediction to seasonal) involving one or more component models (atmosphere, chemistry/aerosols, land, ocean, waves, sea ice, and space weather) that exchange information through the NEMS mediator (often called the coupler). Additionally, at the request of the NCEP director, the UCAR Community Advisory Committee for NCEP (UCACN; <https://vsp.ucar.edu/ucacn>) has formed a Modeling Advisory subcommittee (UMAC; www.earthsystemcog.org/projects/umac_model_advisory/) to review the major modeling components of the production suite along with the near- and longer-term development plans.

The landscape for development of NCEP tools is complex. Some are developed primarily within NCEP (e.g., NMMB), others primarily within NOAA research laboratories (e.g., GFDL hurricane model), others primarily by NCAR (e.g., the ARW dynamic core), and some applications rely on contributions from a variety of sources (e.g., the physics suite used in the RAP model). While NCEP has an ongoing collaboration with the research community in NWP development, there is now increased emphasis on this partnership. The leadership of the NCEP Environmental Modeling Center (EMC) has stated as a priority “making ALL operational codes available with the proper support to make community modeling possible” (Tolman 2015), and this priority is reflected in the draft NGGPS Implementation Plan. The Developmental Testbed Center (DTC) is one of the pillars of the bridge between NWP research and operations. Over the years, the DTC has put in place several mechanisms to facilitate the use

of operational models by the general community, mostly by supporting NWP codes and organizing workshops and tutorials. Additionally, the DTC has established the Mesoscale Modeling Evaluation Testbed (MMET), a set of tools and datasets to enable researchers to demonstrate the merit of their innovations over existing operational baselines. In spite of the relative success of the DTC, and of the transition of several NWP innovations to NCEP operations through these mechanisms, there are

still significant challenges to an effective and efficient collaboration between the research and operational groups, in part driven by the disparity of their missions.

The DTC Science Advisory Board (SAB) has documented the existence of remaining impediments to the use of NCEP operational models in the research community, and recommended continued investment in a software environment that facilitates supporting those models for the community. Similar conclusions were reached by the UCACN, which explicitly recommended NOAA should devote resources to the creation of a modeling infrastructure to facilitate the use of operational suites by the research community. Given this demand, NOAA funded the DTC to design an NWP Information Technology Environment (NITE) to facilitate collaborations and accelerate the transition of innovations from research to NOAA operational modeling suites. This infrastructure is particularly important to help address the UMAC’s recommendation of reducing the complexity of the NCEP production suite through an evidence-driven approach toward decision making and end-to-end modeling development.

NITE DESIGN GOALS. A survey of 40 scientists who run NCEP models was conducted to provide information on the challenges faced by research and development users. The most common issues raised were lack of a friendly way to configure the modeling suites, access limitations to source code repositories, absence of a method to systematically track and reproduce experiments, shortcomings of automation tools, difficulty accessing input datasets, and sparseness of documentation and training.

NITE is envisioned to address these deficiencies, facilitate and enhance the development and transition of innovations to NCEP operational models, and increase the effectiveness of the collaboration between EMC staff and the broader community. In addition to the survey results, its design benefitted from an assessment of selected

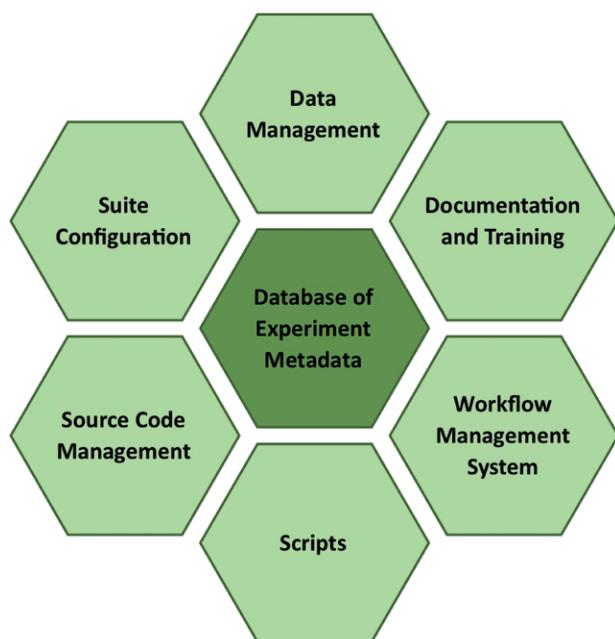


FIG. 3. Infrastructure elements included in NITE: data management, source code management and build systems, suite definition and configuration tools, scripts, workflow management system, experiment database, and documentation and training.

existing NWP infrastructure systems, which was done through review of available documentation; site visits to the European Centre for Medium-Range Weather Forecasts (ECMWF), the Met Office, and the NCAR Community Earth System Model (CESM) group; and focus groups with NWP teams at EMC and the NOAA Earth System Research Laboratory (ESRL). Readers are referred to www.dtcenter.org/eval/NITE for detailed survey results and a description of selected aspects of the infrastructures examined.

ELEMENTS OF THE PROPOSED NITE DESIGN. To fully support model use and development, NITE should have the seven elements described in Fig. 3 and detailed below.

Data management. Typical NWP experiments require a variety of input datasets and produce numerous output files. With the continuous increase in computer power, and the consequent ability to run models at higher resolutions and/or in ensemble mode, the volume of input and output data is ballooning and can only be expected to keep growing. We foresee the need for straightforward archival, browsing, and selective recovery of all inputs and outputs relative to NWP experiments.

All analyses, forecasts, and observations needed as inputs to run the NWP suites to be supported within NITE must be readily available. To avoid loss of information, it is preferable that most datasets be stored in their raw or native form, and converted to other formats when ready to be used. In addition, standardization and publication of data formats and metadata is of paramount importance to facilitate use and assure that all workflow components (such as data assimilation and visualization) can ingest the datasets.

Accessibility of datasets is arguably the most challenging aspect for NITE. The datasets served through the NOAA National Operational Model Archive and Distribution System (NOMADS) are insufficient for initializing experiments with NCEP models. NOAA has an archive that contains most of the input datasets needed for running operational models, but access to external collaborators is limited. While access to proprietary datasets likely cannot be arranged, it is recommended that NCEP make publicly available comprehensive nonrestricted retrospective datasets for select periods of interest plus select challenging forecast cases (the latter is already done through MMET, which could be expanded through cloud-based resources).

Source code management and build systems. Scientists should have access to software management tools to obtain the source code for all workflow components and keep track of their development. A primary goal is that all development should have a path toward potential operational implementation. Code integration with code repositories, such as Git, is a lot easier if all development is synchronized with a single master repository. Scientists should be strongly encouraged to use developmental branches of the code repository to track and recover code, as well as to prevent code from aging off or diverging from the main development. It is not necessary that all code used in a suite be housed in the same code repository, or even at the same institution, as NITE could easily pull code from various places to build all executables

needed for a suite. As an example, NITE would be able to pull in community codes whose authoritative repository is housed outside of NCEP, such as the WRF model housed at NCAR.

Suite definition and configuration. A suite is defined as a set of workflow components that is run using a certain configuration. One example (simplified here for conciseness) is the operational NAM suite, composed of the following workflow components: preprocessor, GSI data assimilation, NMMB model, and the Unified Post Processor (UPP). Likewise, the HWRF suite is composed of preprocessor, GSI, vortex relocation, WRF model, coupler, ocean model, and UPP. It can readily be seen that a single workflow component (such as UPP) can be part of multiple suites.

The NITE design employs the concept of a predefined suite, which is similar to the *component sets* used by the NEMS and CESM communities as described in the CESM v1.2 Users Guide available at www.cesm.ucar.edu/models/cesm1.2/cesm/doc/usersguide/ug.pdf. A predefined suite would be supported out of the box because of its relevance to a group of scientists, such as the operational GFS or a research configuration of HWRF that is being used to support a given field experiment. Scientists could use a predefined suite, whose configuration is stored in the experiment database, as a starting point for conducting their runs or developing their own suite. Or they could completely bypass the predefined suites and build their own suites from scratch.

Component execution scripts. In addition to executing the workflow components, a myriad of small tasks need to be accomplished when running a suite—for example, creating working directories, staging the input datasets, etc. Currently, each NCEP suite has many scripts to accomplish these tasks and there is limited standardization of the scripting systems used by the various suites. Since various suites use the same workflow components, we recommend that the method for calling the components be made uniform, with most aspects of the customization abstracted to configuration files. This standardization will lead to reduced cost of maintenance and training.

One important requirement to ensure portability is that the scripts do not contain any platform-specific or automation features. All automation should be controlled by the workflow management system, described in the next section. This separation of automation from the suite control

scripts would allow the scripts to be used with a variety of automation systems (or none at all).

Workflow management system. Since NCEP suites are complex and contain many tasks, it is generally not feasible for scientists to conduct experiments by submitting jobs individually. This situation is exacerbated for suites representing ensemble prediction systems and coupled model applications. For this reason, users must employ a workflow management system—a software system to manage complex collections of tasks that need to be carried out in a certain way, with elaborate interdependencies and requirements, and to provide fault tolerance.

At NCEP, the ecFlow workflow management system (developed by ECMWF and described at <https://software.ecmwf.int/wiki/display/ECFLOW/ECFLOW+Introduction>) is used to run the operational suites. Since ecFlow is not available on the NOAA research platforms, ESRL has developed the Rocoto workflow management system described at <https://github.com/christopherwharrop/rocoto/wiki/Documentation>), which is portable and can be used on generic platforms. By designing scripts that do not contain automation features, the suites can be constructed to run in operations with ecFlow, in research mode with Rocoto, or even with other workflow management systems available to users.

Experiment database. A database for storage and retrieval of experiment metadata is envisioned for all experiments conducted with NITE. The goal is to record provenance of codes, scripts, configuration files, and inputs related to an experiment so that the experiment can be reviewed and reproduced. This will substantially enhance the level of confidence that EMC staff has in experiments conducted by outside partners.

Documentation and training. Documentation and training on both the NWP systems (suites and workflow components) and NITE itself are of paramount importance for the success of this effort. Since NITE will involve a plethora of documentation, early consideration should be given to standardization of a documentation system that allows for continuous updating.

SUMMARY AND RECOMMENDATIONS TOWARD NITE IMPLEMENTATION. The NCEP modeling suites have become progressively more complex over the years, posing a challenge to

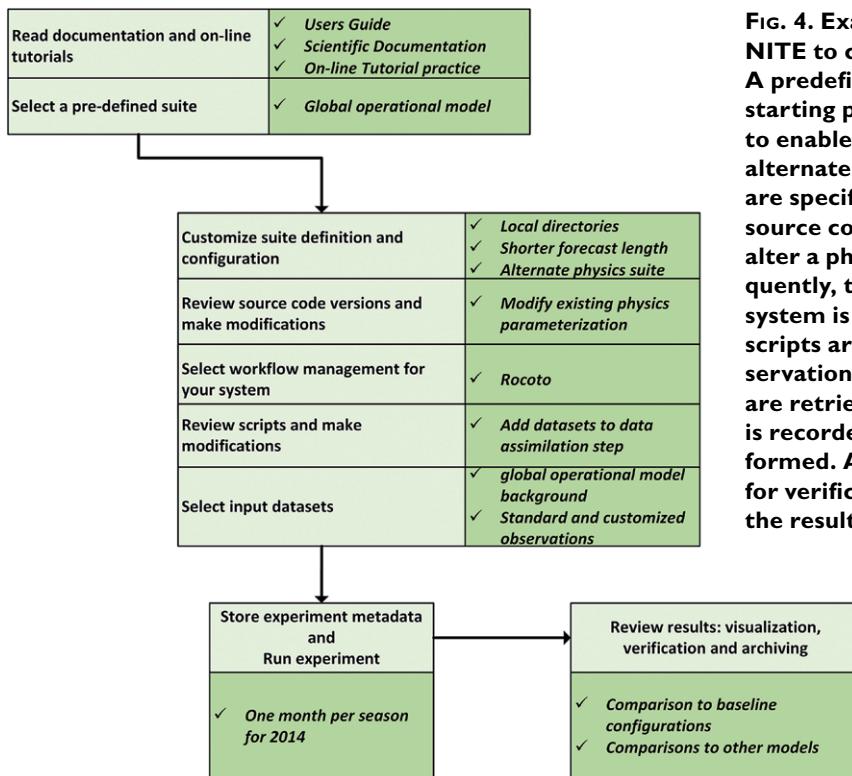


FIG. 4. Example of how a scientist might use NITE to conduct an experiment. A predefined suite, NGGPS, is chosen as a starting point. Suite modifications are made to enable a shorter forecast length and use alternate physics; also, local directories are specified for outputting results. Next, source code is retrieved and modified to alter a physical parameterization. Subsequently, the Rocoto workflow management system is chosen for automation, and the scripts are modified to use nonstandard observations in GSI. Finally, all input datasets are retrieved and staged, the configuration is recorded, and the experiment is performed. Additional workflow components for verification and validation are run, and the results are archived.

scientists inside and outside EMC to configure, launch, and track experiments. This contribution outlined the design for an environment intended to facilitate model development and testing, with the goal of expediting the transition of NWP research to operations.

Several elements are important in this environment, and Fig. 4 shows an example of how a scientist might employ them to conduct an experiment. The design is supportive enough to provide the scientist with a predefined suite as a starting point, yet flexible enough to allow modification of the source code and scripts.

NITE can be implemented incrementally, with the phases planned to best match NOAA strategic goals and funding. For example, there could be an initial deployment on NOAA research platforms to be extended later to community platforms, such as the NCAR supercomputer.

In addition to the MMET developed by DTC, NOAA already has several tools that can be used as a starting point for NITE. Among other examples, the Rocoto workflow management system can be used for automating jobs in nonoperational settings, the HWRF object-oriented Python system codeveloped by EMC and DTC can be used as a prototype for the scripts, the NOAA Virtual Laboratory (VLab) can be leveraged for sharing codes, the experiment database

put forth by NCEP Central Operations (www.nco.ncep.noaa.gov/pmb/docs/Implementation%20Standards%20v10.0.pdf) serve to standardize several aspects of the operational process. Elements of the NITE design and the existing infrastructure within NOAA are being reviewed by the recently formed NGGPS Overarching System Team, which is charged with leading the implementation of NGGPS infrastructure to facilitate NWP development in connection with the community. In fact, EMC and NGGPS are currently engaged in the NITE-based refactoring of the scripts for NCEP's global model, development of a database to capture the provenance information of experiments conducted, development of comprehensive documentation using the DOxygen software, and a process to move some of the source codes to VLab, where they can be more readily accessed.

The NITE development will require a comprehensive effort with significant initial investment, continuous growth, and training, with the expected outcome of extensively facilitating how researchers interact with NCEP operational models. Finally, the NITE code management and experiment database will provide accurate recording of the provenance of experimental results, making them relevant for

used in the WRF Portal can be considered to capture more complete experiment metadata, and the operational implementation standards

consideration by EMC, thus enhancing the number of developments made available for potential operational implementation.

ACKNOWLEDGMENTS. The authors sincerely thank the scientists that participated in the survey and focus groups. We are indebted to Bonny Strong, Stanley Benjamin, and four anonymous reviewers for their constructive comments on an earlier version of this manuscript. Additionally, we thank our colleagues at NCAR, ESRL, EMC, ECMWF, and the Met Office for the valuable input. The DTC is funded by NOAA, the Air Force, the National Science Foundation, and NCAR.

FOR FURTHER READING

- Bernardet, L. and L. Carson, 2015: NITE—Numerical weather prediction information technology environment. Final Report. Developmental Testbed Center, 27 pp. [Available online at www.dtcenter.org/eval/NITE/NITE-report-AOP2014.pdf.]
- , and Coauthors, 2008: The Developmental Testbed Center and its Winter Forecasting Experiment. *Bull. Amer. Meteor. Soc.*, **89**, 611–627, doi:10.1175/BAMS-89-5-611.
- , and Coauthors, 2015: Community support and transition of research to operations for the Hurricane Weather Research and Forecasting model. *Bull. Amer. Meteor. Soc.*, **96**, 953–960, doi:10.1175/BAMS-D-13-00093.1.
- DTC SAB, 2014: Developmental Testbed Center Science Advisory Board meeting report and summary of recommendations. 12 pp. [Available online at www.dtcenter.org/SAB/SABAnnualReport2014_Final_3Nov2014.pdf.]
- Govett, M. W., and J. S. Smith, 2007: On the use of WRF Portal to support the Developmental Testbed Center. *23rd Conf. Interact. Info. Process. Syst.*, San Antonio, TX, Amer. Meteor. Soc., P1.14. [Available online at https://ams.confex.com/ams/87ANNUAL/techprogram/paper_119120.htm.]
- Harrop, C., 2015: Rocoto Workflow Management System. [Available online at <https://github.com/christopherwharrop/rocoto>.]
- Iredell, M., T. Black, and W. Lapenta, 2014: The NOAA Environmental Modeling System at NCEP. *26th Conf. Weather Anal. Forecast./22nd Conf. Numer. Weather Predict.*, Atlanta, GA, Amer. Meteor. Soc., 1B.2. [Available online at <https://ams.confex.com/ams/94Annual/videogateway.cgi/id/26279?recordingid=26279>.]
- Lapenta, W., 2015: The first holistic technical suite review of the Natl. Centers for Environ. Pred. (NCEP) production suite: Why now and how? *27th Conf. Weather Anal. Forecast./23rd Conf. Numer. Weather Predict.*, Chicago, IL, Amer. Meteor. Soc., 1A.1. [Available online at <https://ams.confex.com/ams/27WAF23NWP/webprogram/Paper273390.html>.]
- Ma, X., P. Fox, C. Tilmes, K. Jacobs, and A. Waple, 2014: Capturing provenance of global change information. *Nat. Climate Change*, **4**, 409–413, doi:10.1038/nclimate2141.
- NOAA, 2014: R2O Initiative: Next Generation Global Prediction System (NGGPS) Implementation Plan. Draft as of October 10, 2014 v 1.0, 51 pp. [Available online at www.nws.noaa.gov/ost/nggps/NGGPS%20Implementation%20Plan%20v1.0.pdf.]
- Smith, S., 2015: Update on NOAA's Virtual Lab. *NOAA Testbed and Proving Ground Workshop*, Boulder, CO. [Available online at www.testbeds.noaa.gov/events/2015/workshop/presentations/tue-smith-vlab.pdf.]
- Tolman, H., 2015: The NCEP production suite: Present and future. *27th Conf. Weather Anal. Forecast./23rd Conf. Numer. Weather Predict.*, Chicago, IL, Amer. Meteor. Soc., 1A.2. [Available online at <https://ams.confex.com/ams/27WAF23NWP/webprogram/Paper273797.html>.]
- Trahan, S. G., and Coauthors, 2015: Modernizing the operational workflow and automation of the NCEP Hurricane Weather Research and Forecast (HWRF) modeling system using Python and Rocoto. *5th AMS Symp. Adv. Model. Anal. Using Python*, San Diego, CA, Amer. Meteor. Soc., 1.3. [Available online at <https://ams.confex.com/ams/95Annual/webprogram/Paper267665.html>.]
- UCAR, 2014: Annual Report of the UCAR Community Advisory Committee for NCEP. 51 pp. [Available online at www.cpaess.ucar.edu/sites/default/files/images/ucacn/UCACN_Report_Jan2014_v10-22Sep2014.pdf.]
- , 2015: Report of the UCACN Model Advisory Committee. 72 pp. [Available online at www.ncep.noaa.gov/director/ucar_reports/ucacn_20151207/UMAC_Final_Report_20151207-v14.pdf.]
- Wolff, J. K., and Coauthors, 2016: Mesoscale Model Evaluation Testbed (MMET): A resource for transitioning NWP innovations from research to operations (R2O). *Bull. Amer. Meteor. Soc.*, **97**, 2135–2147, doi:10.1175/BAMS-D-15-00001.1.