System Architecture for Operational Needs and Research Collaborations

March 31, 2017

Initial report of the NOAA NCEP System Architecture Working Group
https://esgf.esrl.noaa.gov/projects/sawg

Introduction

This is the first report of the System Architecture Working Group (SAWG), formed in October 2016 by Mike Farrar, Director, Environmental Modeling Center (EMC), National Centers for Environmental Prediction (NCEP), National Weather Service (NWS), National Oceanic and Atmospheric Administration (NOAA). The SAWG was charged with providing prioritized recommendations on the system architecture for operational prediction system employed by NCEP to produce numerical guidance for operational weather, climate and environmental predictions. The system architecture also serves the purpose of enabling and encouraging collaboration with organizations and individuals within the research community (see Appendix A for the full charge and Appendix B for a list of the SAWG members). This initial report was developed in a series of conference calls during the period 21 October 2016 - 10 March 2017.

We begin with a brief discussion of what the system architecture encompasses. The central part of the report is a list of recommendations, organized into several categories and prioritized as essential or desirable. The *general*, *structural*, and *technical* recommendations relate to the entire system architecture. They are based on background material that includes a system architecture description and glossary[1], additional definitions from the SAWG - see Appendix F, and recommendations from the UCACN Model Advisory Committee (UMAC), a sub-committee of the University Corporation for Atmospheric Research (UCAR) Community Advisory Committee for NCEP (UCACN) – see Appendix G.

Recommendations for *modeling applications* relate mainly to the NOAA Environmental Modeling System, or NEMS. NEMS is the part of the system architecture that supports the coupling of model components into multiple applications, for different predictive targets (e.g. space weather, seasonal). It is an initial focus because of its central role in the construction of a unified modeling system that spans these applications, and in supporting interoperability of components with other national centers. These recommendations combine the general, structural, and technical recommendations with evidence in the form of case studies of component interoperability (Appendix C), a description of the requirements for NEMS-based, coupled applications and pointers to documentation of delivered milestones (Appendix D), and a NEMS gap analysis (Appendix E). We anticipate expanding the focus of the SAWG to other critically important aspects of the system architecture in a final report. In particular, we expect to address aspects of workflows, data assimilation, libraries and utilities, and the atmospheric physics interface, including issues pertaining to atmospheric chemistry and aerosols, that relate to the system architecture.

Prior to formulating priorities, a set of questions to be addressed by the system architecture was developed by the SAWG[2].

---

[1] Auligne et al. 2016, NEMS System Architecture Overview, see https://esgf.esrl.noaa.gov/site_media/projects/nems-workshop/report_1610_system_architecture.docx

[2] https://tinyurl.com/sawg-questions

Vision

Operational prediction employs the best possible numerical forecast system, embodying all relevant and evidence-based knowledge at any given time. Collaborative research employing the operational prediction system or its components is best situated to realize this vision.

System Architecture at NCEP

System architecture can be defined as "the fundamental organization of a system, embodied in its components, their relationships to each other and the environment, and the principles that govern its design and evolution."[3,4] The software system architecture used by the National Weather Service (NWS) at NCEP/EMC is critical because it serves as the backbone of a unified modeling system, and must provide high performance, reliable technical and scientific functions for a range of different forecast products. The design of the architecture is relevant to research community partners because it must make it easy for them to perform runs and experiments, and participate as full partners in model development. Here we consider the national "community" to be comprised of government labs, universities, and other organizations with an interest in developing and using the coupled prediction system for research. A more complete and detailed definition of the community is left to the April 2017 NOAA Community Modeling Workshop.

A shared view of the scope and elements of the unified modeling system architecture at NCEP/EMC emerged from a workshop held there on September 1-2, 2016.[5] At that workshop, about 90 participants formed teams and developed diagrams of the system architecture. These were consolidated into a single diagram (Figure 1) and described in a short document and glossary. This system architecture is a layered, component-based structure, divided into (1) a Workflow Environment that includes a user interface and database of experiment metadata for previous runs, including metadata about input datasets and observations/analyses used for verification, (2) a Prediction Package layer that consists of a sequence of pre-processing, data assimilation, forecast, and post-processing jobs, (3) a Modeling and Data Assimilation Application layer that includes the coupling framework (the NOAA Environmental Modeling System, or NEMS), a prescribed interface between atmospheric physics and dynamics, model components, and data assimilation components, and (4) a layer of Libraries and Utilities. Each layer utilizes components, which can be defined as "composable" software elements that have a clear function and interface. The system architecture includes elements that are complete and others that are still in progress.

---

[3] See for example IEEE/ISO/IEC 42010-2011.
[4] The system architecture should be distinguished from the *software infrastructure*. The software infrastructure is a set of technical building blocks that represent a wide range of implementation options. The system architecture defines what choices are made and what is built; the software infrastructure is a set of tools for building it.
[5] NEMS Code, Data, and Documentation Management Workshop, College Park, MD, Sept. 1-2, 2016, https://esgf.esrl.noaa.gov/projects/nems-workshop/

The portion of the system diagram that relates to coupled modeling applications is shown in teal and black. NEMS is shown in teal and includes a main coupler, a space weather coupler, a driver, and tools for building applications and running specific cases.

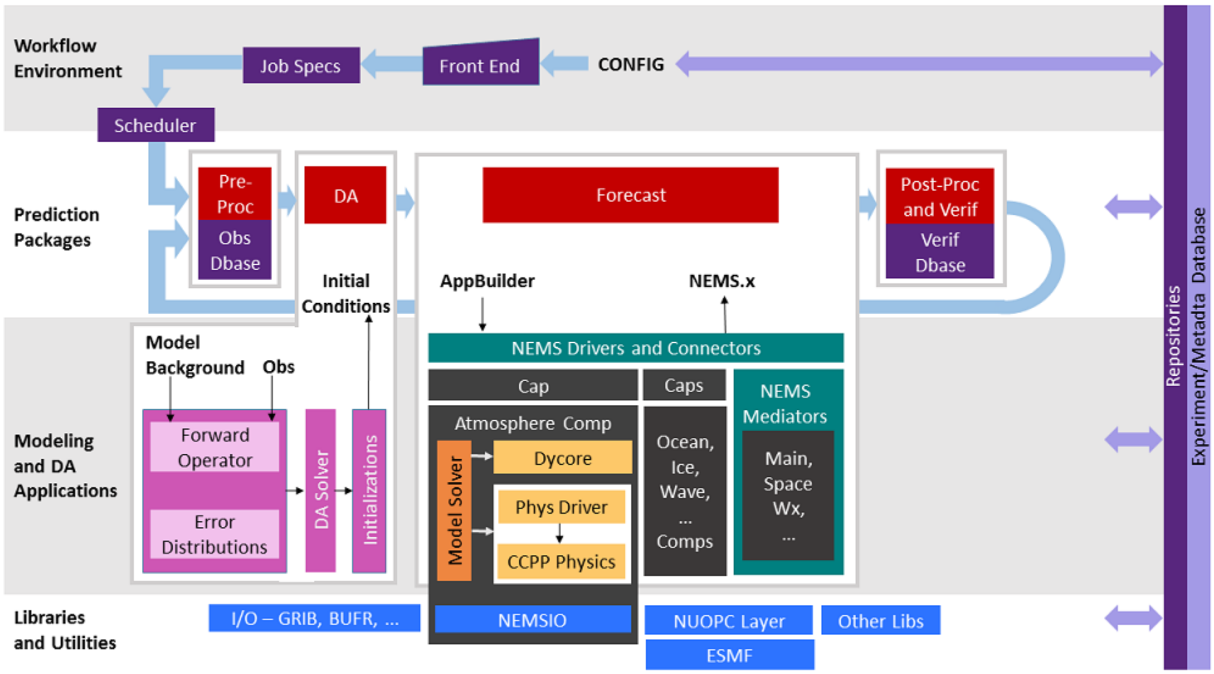

Figure 1. Diagram showing the four main layers in the unified modeling system architecture: Libraries and Utilities, Modeling and Data Assimilation Applications, Prediction Packages, and Workflow Environment. Purple boxes indicate parts of the Workflow Environment and databases, with thick light blue lines indicating sequence. Red boxes indicate executables while the thin lines around them represent scripts that invoke the executables. Teal boxes show NEMS infrastructure. Black boxes represent science components, caps, and mediator components. Orange boxes show subcomponents of the atmosphere model component. Pink boxes show parts of the data assimilation system. Blue boxes show utilities and libraries. The Prediction Package sequence shown is typical; it may change for different applications.

Recommendations

The recommendations listed below are grouped in categories: General, Structural, Technical and specifics for Modeling Applications. We have also had preliminary discussions of recommendations for data assimilation applications; however, the SAWG has not had sufficient time to fully explore these recommendations, so they are included as an appendix and will be taken up in a future SAWG report. We recognize that it will take some time and effort to implement our recommendations, which are divided into those that are "desirable" and "essential". Essential recommendations need not be implemented immediately, but we suggest they must be considered and included in the project from the start, to avoid choices during planning and development that would preclude their eventual adherence, or require costly, difficult and lengthy revisions to the software architecture as the project evolves. We hope recommendations listed as desirable are also followed, but we think that it is less critical that they drive development activities. Furthermore, we note that it is possible that two "essential" recommendations may be in conflict at some stage of implementation and we have not offered a way to resolve that conflict. Also, we realize that some of these recommendations may be more actionable than others. We will take up these matters in future SAWG discussions. Finally, it is important to point out that the recommendations are not ranked in order of importance.

It should be noted that members of the SAWG have had several discussions of process, including how decisions are made with respect to the system architecture and how to engage partner organizations in advancing software codes and practices. While these issues are to some degree beyond the scope of the SAWG's charge and more appropriately the purview of the Governance working group, we include governance-related recommendations that we believe are necessary to ensure that the system architecture operates efficiently, evolves to meet new requirements, ports to other platforms, and is interoperable with infrastructure systems in place at partner institutions.

General Recommendations

	Importance	**Recommendation**
G1	Essential	Meet the needs of stakeholders (the ultimate judges of architectural quality), including entities and individuals within the operational forecast and research communities.
G2	Essential	Be cost effective and timely.
G3	Essential	Allow for maintenance, iterative design, evolution, extensibility, etc., as required by both operational and research stakeholders. Ensure that interfaces are clear with a formal procedure for extension.
G4	Essential	Employ modularity and object-oriented design in order to enable scientists to focus on the portion of code of interest to them.
G5	Essential	Balance local (within component) vs. global optimality.
G6	Essential	Acknowledge, manage, and mitigate risks and be able to continue to function if risks materialize. Risks include hardware and software failures, inability of the architecture to meet requirements, and inability of the architecture to adapt to changing scientific or computational needs.
G7	Essential	Adopt standards to the extent possible (coding, component interface, etc.).
G8	Essential	Be as interoperable as possible with architectures (e.g. coupler configurations) in U.S. partner institutions.
G9	Essential	Software should have technical, scientific, and user documentation that is posted in a publically accessible location.
G10	Essential	Use modern development tools and processes to improve ease and quality of development, and foster collaboration between organizations. Examples are ticketing systems, the cohesive use of version control software, integrated development environments, and agile software development.
G11	Desirable	Maximize the value of the system by focusing on external interfaces, form, and delivered function.
G12	Desirable	Limit complexity to that which is essential, i.e., required for robust functionality.

Structural Recommendations

	Importance	Recommendation
S1	Essential	Be organized in a series of layers, including libraries, modeling applications, prediction suites, and workflow (Fig. 1), in a manner consistent with operational standards defined in the Environmental Equivalence 2 (EE2) document.[6]
S2	Essential	Support application development by independent groups according to their own timelines, while also sharing components and infrastructure as part of a unified modeling system. In order to continue with unified regression testing during branch phase, components retain a backward compatible option when new features or fixes are implemented. This needs to be paired with a governance process that limits divergence of independent development paths, enables pruning regression and scientific test requirements, and reviews code and processes for obsolescence.
S3	Essential	Address the requirements of an initial list of applications and components (in each case, “prediction” is assumed): Weather, Sub-seasonal, Seasonal, Ocean Surface Waves, Sea Ice, Whole Atmosphere (Stratosphere, Mesosphere, Ionosphere, Plasmasphere, and Electrodynamics), Regional Land-surface Hydrology, Regional Nest, Air Quality, and Coastal, provided by science leads. The applications should include a definition of which Earth system components are prognostic (“active”) for a given application.
S4	Essential	Be able to support explicit coupling among system components, and also have the flexibility to implement implicit coupling between pairs of components as required.
S5	Essential	Have a scripting/workflow infrastructure, along with clear application programing interfaces, that support deployment of the modeling and DA applications at multiple organizations. Implementing this infrastructure in a common multi-capability language, like Python, would enable a wider community to engage in its development. Convergence of scripting/workflow infrastructure between research and operations would also streamline research to operations.
S6	Essential	Support both prognostic and non-prognostic versions of components to enable testing, development, and mechanism-withholding experiments. Examples of non-prognostic components are prescribed data components and no-op components.
S7	Desirable	Support component hierarchies, which allow components to drive other components. This allows for code encapsulation of sub-processes.
S8	Desirable	Support a rich set of ensemble construction arrangements, including multiple instantiations of individual components (either averaged or individually selected), multiple instantiations of sub-grid scale parameterizations, and full communication among members either within a single component or among multiple components.

[6] https://esgf.esrl.noaa.gov/site_media/projects/nems-workshop/EE2Structurefinaldraft.docx

Technical Recommendations

	Importance	Recommendation
T1	Essential	Publicly document common requirements for the interface between atmospheric dynamics and sub-grid scale (SGS) parameterizations, including (possibly different) staggers, grids, distributions, time loops, or processor sets. Include potential to evolve to three-dimensional treatment of sub-grid scale parameterizations.
T2	Essential	Publicly document coupling requirements, including requirements for components to run on the same grid and in a particular order, computation of fluxes between pairs of components (e.g. exchange grid), and interpolation methods.
T3	Essential	Publicly document common and application-specific requirements for history file fields, periods, averaging, multiple streams and control.
T4	Essential	Publicly document common and application-specific requirements for ensembles and data assimilation. Note: The SAWG has begun to formulate recommendations for data assimilation applications, but has not reached a consensus. A preliminary set of recommendations is provided in Appendix H as an indication of the tenor of the discussion that is anticipated.
T5	Essential	Publicly document common and application-specific documents for workflows and scripting, including any scripting for configuration, regression testing and restarts.
T6	Essential	Support diagnostic interrogation of model output for testing (e.g. experimental alternative physics packages), model evaluation (e.g. standard set of re-forecast metrics), and operational prediction quality assessment. This may include some in-line diagnostic computation. This should include enough of the operational products and verification methods to determine whether a research development improves the forecast.
T7	Essential	Enable high scalability to ensure optimal time to solution and total cost of ownership on current and emerging large, high-performance computer systems such as are employed operationally and in large research installations.
T8	Essential	Be aware of emerging software and data structure constraints, as HPC exascale computing and data transfer needs evolve. Attempt to avoid design choices that might impair adaptability to those constraints.
T9	Desirable	Support interchangeability of file formats for both internal use and external input/output to and from the workflow. See NGGPS-CDDM document[7].

[7] Code, Data, and Document, Management for NEMS Modeling Applications and Suites, prepared by the NGGPS Overarching System (OAS) team, https://tinyurl.com/nems-cddm

Recommendations for Modeling Applications

	Importance	Synthesis Recommendation - Modeling Applications
MA1	Essential	The SAWG proposes initiating discussions and formalizing tests that explore the feasibility of the following combination: ● Base the seasonal EMC scientific choices on an existing modeling system, and transition coupled system development to a body in which scientific leadership is shared with external participants. ● Utilize NEMS or NEMS-like ESMF/NUOPC-based infrastructure at EMC in order to facilitate interoperability with partners and leverage shared tools and technical support.
MA2	Essential	Understand current best practices and restructure build and other scripts.
MA3	Essential	Explore partnership with the CESM project to define ways to engage coupled system science contributors from the broader community, develop community-friendly infrastructure, and leverage established outreach and training programs.
MA4	Essential	Establish a standing science lead or science steering committee responsible for the direction of the overall NOAA unified modeling system.
MA5	Essential	Name a modeling system lead at EMC who can serve as the primary point of contact and coordinator for coupling science and technology.
MA6	Essential	In collaboration with NCO, establish formal processes at EMC that allow for external participation in technical and scientific decision-making.
MA7	Desirable	NOAA, NCAR and partner organizations work toward supporting the NEMS or a similar coupler as community software.

Discussion - Modeling Applications

This section is intended to synthesize recommendations from the previous sections and collected evidence. The discussion here is focused on the development of coupled modeling applications that are part of the unified modeling system. Where a recommendation is relevant to the discussion, its number is indicated. Not all recommendations are discussed.

Shift to Community Development

Several critical components of the NOAA unified modeling system are shifting to development teams outside of NCEP EMC. For example, the new atmospheric dynamical core will come from the Geophysical Fluid Dynamics Laboratory (GFDL). Other components in the unified modeling system, including ocean models and ice models, are likely to originate at other centers.

The SAWG recommends extending this approach to include the coupling science and implementation choices leading to the final predictive applications. The rationale is that community engagement is critical to address the complexity of coupled model development. The advancement of a suite of coupled applications will require a highly coordinated organization, significant staff support, expertise in both coupling science and the constituent components of the modeling systems, and carefully considered governance. For these reasons, most coupled model development is best undertaken outside of EMC, in a manner that supports both focused development toward operational use and community contributions. This kind of engagement has already begun. The SAWG notes that there are positive developments in this direction: pilot activities for EMC collaboration with GFDL on the seasonal and sub-seasonal coupled system, with CESM on community support for coupling infrastructure, and with Navy and NASA on integration of specific components such as ocean waves and chemistry, respectively, in addition to long-standing collaboration on data assimilation.

Coupling Infrastructure

Ideally, the underlying software infrastructure and system architecture will provide a means for the community to collaborate on code development of the unified modeling system, and to adapt to future technical and scientific challenges. Whether or not to continue using the NOAA Environmental Modeling System (NEMS) is a central question for the system architecture going forward. The NEMS infrastructure was initially developed by EMC staff to create a shared framework for the mesoscale and global atmospheric prediction component models. It was updated and extended by the ESMF technical team and others over the last several years, with guidance on coupling science decisions provided mainly by EMC scientists.

During this time, the NEMS architecture evolved to support a variety of prototype applications, including weather prediction, seasonal prediction, a regional-nest application, land/hydrologic processes, and coupled space weather (S3). The source of requirements for these applications is described in Appendix D. The diversity and nature of these applications is an important consideration because they reflect requirements for a range of coupling techniques, highly flexible grid remapping software, flexible run sequences, and technical interoperability with components originating from NOAA Geophysical Fluid Dynamics Laboratory (GFDL), the NOAA Space Weather Prediction Center (SWPC), the Naval Research Laboratory (NRL), the National Center for Atmospheric Research (NCAR), NASA Goddard, and other organizations (G1, G8).

Under the NEMS infrastructure, these applications are handled in a unified manner so that they have a consistent checkout, build, and run capability, and can share coupling tools and components systematically. At the same time, NEMS offers a formalized “app” structure that enables application development groups to work independently, and proceed on different timelines (S2). This architecture is one of the unique features of NEMS that addresses the architectural diversity of the components that it supports. NEMS can be embedded in a layered system architecture, as shown in Figure 1 (S1).

The component and coupling aspects of NEMS are based on ESMF/NUOPC. ESMF is a general, multi-agency framework that is used by NOAA partners including NASA, the Navy, and CESM, along with thousands of individual users. NUOPC Layer interfaces, which are bundled in the same distribution, were introduced in 2011. The NUOPC code is mature and adding features, and has been used to develop a range of coupled applications. Theurich et al. 2016 describes the ESMF and NUOPC Layer adoption status at NASA, the Navy, CESM, and NOAA. This user base and support level represent a practical advantage in that ESMF/NUOPC code has been extended to address a wide variety of research and operational problems and applications, and key capabilities are heavily vetted. The establishment of ESMF/NUOPC as a component interface standard and community modeling framework represents an advancement toward increasingly interoperable coupled modeling systems at federal centers, and many years of investment by the U.S. modeling agencies (G7). A set of interoperability case studies with components relevant to NEMS is included as Appendix C. In addition to the coupling of major modeling components (atmosphere, ocean, ice, etc.) which is the focus of NUOPC, coupling issues pertaining to atmospheric chemistry and aerosols, and its tight connection to atmospheric physics, is a topic necessitating further analysis and consideration.

Science Challenges

Many model components have been run under NEMS and validated bit for bit, with few changes to the underlying user code, and the initial coupled systems demonstrate feasibility of the general approach. However, it has been difficult to transition from the initial coupled prototypes to models with increasing predictive skill. The integration efforts within EMC lack staff, governance, and organizational mandates compared to the teams dedicated to the integration of coupled systems at centers like GFDL and NCAR. An ideal solution would engage external expertise to accelerate scientific improvements, would leverage community infrastructure and encourage collaboration, and would address the specialized application and operational needs of the NWS.

Alternatives, Opportunities, and Costs

GFDL is a natural place to look for both coupling science expertise and coupling infrastructure. Several of the components of the seasonal system are expected to come from GFDL, which maintains an institutional infrastructure called the Flexible Modeling System (FMS). An advantage of FMS is that it has been used to develop high-quality coupled models that include several of the key components desired in the NOAA unified modeling system. This model development represents many years of investment of GFDL science and technical resources. A disadvantage of FMS is that it is not a general, established community framework, with a governance structure that allows for community contributions, and a user support and outreach team dedicated to the infrastructure.

An important realization is that the implementation or replication of specific coupling strategies represents a far smaller investment. One way to measure this is through lines of code. While ESMF represents about a million lines of code, a typical coupler built using ESMF tools is much smaller, at thousands of lines of code. Individual couplers are effectively customizations of a large shared software base. Interoperability is supported, even with different couplers, through the underlying standard component interfaces. Some key differences in scientific approach at different centers, such as using an exchange grid (e.g. GFDL and NASA) or not (e.g. CESM) can be implemented as switches. Changes to the NEMS and similar couplers to implement or replicate specific science options are not expected to be onerous.

Avoiding the creation of a new coupled system or a new community modeling framework, both major endeavors, will contribute to goals of cost effectiveness and timeliness (G2). The question of whether or not to continue with NEMS, and not just ESMF/NUOPC, was approached through an initial gap analysis - see Appendix E. This analysis indicates opportunities for improvement, but no showstopping technical issues.

The SAWG recommendation is that the coupled, ESMF-based NEMS system should be used to replicate an existing set of science choices, such as those represented in a coupled GFDL model (MA1). This approach supports the desire for expedience in getting high-quality coupled systems running, and will leverage established community infrastructure. It can be implemented within the "app" based NEMS architecture that has proven flexible enough to support a diverse set of component arrangements and coupling strategies. A potential disadvantage to this approach is inheriting legacy code and scripts that have proven difficult to maintain, understand, extend, and use. In general, these are not part of NEMS but are used for model initialization, data ingest, and execution. *The SAWG recommends an active effort to understand current best practices and refactor these scripts* (MA2).

The solution described above still lacks a mechanism for community engagement in coupled system development. For this we look to the CESM project, whose model supports several of the components in use at EMC and GFDL (MA3). Soon, it will include the FV3 dynamical core in its main development code, and will also be integrating GFDL's MOM6 as its primary ocean model. There are current pilot projects that support development of a next-generation CESM coupler based on the ESMF/NUOPC infrastructure, in collaboration with EMC and the Navy. This presents exciting opportunities for GFDL, CESM, and the NWS to work together with the broader community to support the science and technology of coupled modeling. *The SAWG recommends that NOAA, NCAR and partner organizations work toward supporting NEMS or similar infrastructure code (coupler, driver, etc.) as community software* (MA7). A key opportunity is to explore the possibility of leveraging established CESM training programs and outreach in support of the EMC unified modeling system.This opportunity will require an understanding of the roles to be adopted by GFDL and CESM, and coordination between the CESM project and the Global Model Test Bed, which is offering a community interface to atmospheric model physics.

Path Forward

To move forward in an evidence-based manner, utilizing all available knowledge, the SAWG recommends proceeding along two paths. The technical path explores the feasibility of replicating an existing science approach using NEMS. The SAWG recommends that NEMS/ESMF, GFDL, CESM, EMC and others work together to define a set of test problems and metrics that explore areas such as: 1) significant differences in framework capabilities; 2) how readily GFDL components can be moved into NEMS, and how readily NEMS components can be moved to the new CESM coupler; and 3) the feasibility of replicating the GFDL coupling approach in NEMS. The rationale for working together is that the development teams have limited knowledge of the other teams' software.

The second path has to do with the evolution of the coupling science for the unified modeling system, and the governance necessary to operate a complex system (S3). The NEMS gap analysis in Appendix E identifies several significant management gaps that affect the design and construction of modeling applications, with the lack of a clear overall governance body for the unified modeling system foremost. *The SAWG recommends establishing a standing lead or steering body to make decisions about the direction of the NOAA unified modeling system overall* (MA4). This entity would be responsible for defining evidence-based strategies for making science choices about the coupling that will support the range of projected applications in the unified modeling system.

Staffing choices at EMC will be critical to the success of this multi-organization partnership. In addition to establishing overall science leadership, *the SAWG sees the urgent need for an overall modeling system lead who can serve as the hands-on coordinator for coupling science and technology at EMC* (MA5). This modeling system lead must be able to communicate and manage processes that involve multiple development partners. In addition, success depends upon the transition at EMC to formalized decision-making that includes external inputs. These processes must be established in concert with the NCO to ensure that operational needs are met.

Appendix A - Charge to System Architecture Working Group

To provide the NCEP/EMC Director with prioritized recommendations for the advancement of a system architecture that meets operational needs as well as enables and encourages collaboration with external model development partners and the broader research community.

The System Architecture Working Group (SAWG) should follow an evidence-driven approach, and to that end may review or request plans, performance tests, requirements, reports, case studies, etc. The SAWG should address the scientific and technical quality of the current implementation, identify causes of delays and difficulties, and propose approaches to resolving difficulties and open questions. Recommendations may extend to the organizational aspects of supporting a coupled, unified modeling system and the development of procedures and policies that may be needed to maintain and operate a community-based system architecture. While this initial high-level charge comes from the EMC Director, the SAWG has the latitude to address any other specific questions they feel practicable and appropriate to the issues at hand. The final recommendations can be delivered in whatever form the SAWG determines (white paper, briefing slide, etc.).

Scope:
The SAWG will be limited to a finite period, and will stand down after the delivery of a prioritized set of final recommendations to the EMC Director. However as part of the final set of recommendations, this limited WG should consider whether a new standing Working Group should be established to advise on SA issues as the architecture continues to evolve.

Appointment of Members:
The co-chairs of the SAWG will be appointed by the Director of NCEP/EMC, after consultation with the Director of NCEP. The other SAWG members will be appointed by consensus of the EMC Director and SAWG co-chairs, and will be drawn from EMC, other NOAA development organizations, and other scientific community members as appropriate.

Meetings:
The primary mode of communication will be a biweekly call. Invitees to the call may include the application leads for NCEP/EMC modeling applications, component liaisons, NGGPS team leads, or other teams as identified by the NCEP/EMC director or SAWG chair(s). Invitees can designate an alternative representative if desired.

Appendix B – SAWG Members

Auligne, Tom	Joint Center for Satellite Data Assimilation (JCSDA)
Balaji, V.	Princeton University
Benson, Rusty	NOAA Geophysical Fluid Dynamics Laboratory (GFDL)
Bernardet, Ligia	NOAA Earth System Research Laboratory (ESRL)
Chawla , Arun	NOAA NCEP
Chu, Philip	NOAA Great Lakes Environmental Research Laboratory (GLERL)
Craig, Tony	National Center for Atmospheric Research (NCAR)
da Silva, Arlindo	NASA Goddard Space Flight Center (GSFC)
DeLuca, Cecelia (co-chair)	NOAA Earth System Research Laboratory
Derber, John	NOAA NCEP
Doyle, Jim	Naval Research Laboratory (NRL)
Farrar, Michael (ex officio)	NOAA NCEP
Iredell, Mark	NOAA NCEP
Kinter, Jim (co-chair)	George Mason University (GMU), Center for Ocean-Land-Atmosphere Studies (COLA)
Lamarque, Jean-Francois	NCAR
Michalakes, John	UCAR/CPAESS
Rasch, Phil	DOE Pacific Northwest National Laboratory (PNNL)
Saha, Suranjana	NOAA NCEP
Tallapragada, Vijay	NOAA NCEP
Theurich, Gerhard	Earth System Modeling Framework (ESMF)
Trahan, Sam	NOAA NCEP
Vertenstein, Mariana	NCAR
Wang, Jun	NOAA NCEP

Appendix C - Interoperability Case Studies

WAVEWATCH III: A NUOPC-compliant WAVEWATCH III component was developed at the Naval Research Laboratory (NRL) by Tim Campbell, as part of a regional coupled modeling system. This component became the starting point for integration of WAVEWATCH III into NEMS. The code and inputs were provided by NRL in a repository at EMC on 11/15/2015. The code was running standalone under NEMS by 12/5/2016. Jessica Meixner started as an EMC wave model developer in January 2016, started running on theia in February, and had a one-way atmosphere to wave coupled system working in March 2016. Bug fixes were shared with the NRL COAMPS team and the developers charged with integrating WAVEWATCH III into the global NRL coupled model. Wave modelers at EMC estimate that the shared standard interface saved about 6 months of work.

HYCOM: A NUOPC-compliant HYbrid Coordinate Ocean Model (HYCOM) component was developed by James Chen at NRL, as part of a global coupled modeling system. The component became the starting point for integration of HYCOM into NEMS. The component was received from NRL on 3/18/2013, ported to Zeus, and placed with appropriate inputs for coupling in an EMC repository on 5/3/2013. The code was validated running standalone under NEMS by 7/12/2013. HYCOM coupling at EMC was set aside as Modular Ocean Model 5 (MOM5) coupling in NEMS took priority. In the meantime, the NUOPC version of HYCOM was coupled to the Community Earth System Model (CESM), a process that fixed several issues with the HYCOM cap. After a coupled MOM5 baseline (UGCS-Seasonal 0.2) was established in April 2016, work picked up again at EMC with coupling HYCOM. A single-domain regional coupled system that used NUOPC HYCOM was delivered on 10/12/2016 and an initial nested coupled version was delivered on 11/18/2016. A global coupled system with HYCOM was also delivered at EMC during early 2017. Sharing component interfaces meant that technical and science advancements from NRL and CESM could be leveraged to accelerate the pace of coupled system development at EMC.

LIS/Noah land and WRF-Hydro: NUOPC caps for the NASA Land Information System (LIS) and Weather Research and Forecast Model-Hydrological model (WRF-Hydro) were developed under National Science Foundation funding, and integrated into NEMS. A first milestone of both components running standalone was completed on 9/1/2015 and a technical coupling exchange milestone was completed in 2/22/2016. Since a separate land component was lower priority at EMC, work shifted to development of a regional coupled system at Navy that used the same NUOPC LIS component coupled to the Coupled Ocean Atmosphere Mesoscale Prediction System (COAMPS) atmosphere, in a nested configuration. This milestone was completed on 2/22/2017. The next milestone for the coupled COAMPS work will include the NUOPC WRF-

Hydro component. At NOAA, work has resumed on implementing a separate land component, leveraging knowledge and code from the LIS-COAMPS project. The next EMC milestone, which is LIS/Noah land running as a separate component in the global coupled seasonal system, is scheduled for April 2017. A NASA project to couple WRF-Hydro and LIS and implement coupled DA is also underway, and it is leveraging the NUOPC LIS and WRF-Hydro interfaces that EMC and Navy have developed. The three agencies that are working with WRF-Hydro and LIS coupling are leveraging code and knowledge directly.

NGAC: NCEP has just implemented operationally v2.2 of the NEMS Global Forecast System (GFS) Aerosol Component (NGAC[8]). This is based on the same GOCART ESMF Grid Component used in NASA's GEOS-5 model. This is another example of a demonstrated interoperability exercise where a non-NEMS component can be integrated into a NEMS environment, for an operational application. Coupling of GOCART Grid component relies on the availability of the Physics Grid Component in this version of NEMS.

Documentation of NEMS milestones is here, under Milestone Revisions on the left navigation bar: https://esgf.esrl.noaa.gov/projects/couplednems/

[8] NGAC Release Notes, https://esgf.esrl.noaa.gov/site_media/projects/sawg/Release_Notes_NGAC.v2.2.0.docx

Appendix D - Sources of Application Requirements

The UCACN recommended that NOAA develop a unified modeling system, defined as one in which applications that span global-to-local domains and weather to climate predictive scales share model components and infrastructure software[9]. The current implementation of this unified modeling system is the NOAA Environmental Modeling System, or NEMS. All NEMS applications currently use ESMF/NUOPC Layer version 7 as a coupling framework.

Requirements for coupled applications come from a variety of sources. The primary requirements for NEMS coupled applications come from a matrix of applications and their constituent components (see: http://tinyurl.com/nems-apps). This matrix includes component integration sequences, timelines, and application leads. It was reviewed and confirmed by the EMC director (Tolman) at creation. The NGGPS implementation plan contains supporting information about desired components and applications.

Information in the application matrix has been updated in consultation with EMC managers and application leads. It is not complete for all applications, especially for applications where coupled configurations were given a lower priority. Targeted requirements collection activities were initiated as needed; see draft documents at: https://tinyurl.com/noaa-apps-reqs). The requirements documents have been difficult to finalize in the absence of a decision-making body and process that spans NEMS applications.

Delivery of each application is structured as a sequence of milestones. Milestone documentation is linked to the application matrix. Milestones have included a three-way coupled atmosphere-ocean-ice application (GSM-MOM-CICE), a one-way coupled atmosphere-ionosphere application (WAM-IPE), a regional nest application (NMMB-HYCOM), a wave-atmosphere application (GSM-WAVEWATCH-III), and an initial technical coupling of separate land and hydrology components. Documentation pages are under Milestone Revisions at https://esgf.esrl.noaa.gov/projects/couplednems/).

Requirements for the underlying ESMF infrastructure were collected in an exhaustive community requirements process, and its development has been guided by more than a decade of user inputs. The NUOPC Layer, introduced in 2011, was developed through a community process led by an interagency Common Model Architecture committee.

[9] Auligne et al. 2016, NEMS System Architecture Description.

Appendix E - NEMS Gap Analysis

As described in Appendix D, NEMS is serving as the modeling application infrastructure of the current NOAA modeling system. NEMS is under active development, and, as such, is an incomplete implementation of the fully envisioned unified modeling system. As part of SAWG's initial deliberations, it was determined that an analysis is needed of the current status of NEMS in relation to a perceived ultimate system architecture. In this appendix, we provide a preliminary gap analysis, but thorough and ongoing evaluation is needed as NEMS and the unified modeling system evolve.

NEMS currently consists of a NEMS driver, main NEMS mediator (coupler), a space weather mediator, configuration files that specify processor layouts and component run sequences, standard case specifications (compsets) and a CompsetRun utility, a regression test system, and an AppBuilder script that builds any of the models that have been integrated into NEMS. NEMS code is available to collaborators but it is not licensed for community use.

NEMS makes use of but does not include: the ESMF/NUOPC community infrastructure, the interoperable physics driver and Common Community Physics Package, and specific model components. The NEMSIO package developed by EMC is not considered part of NEMS by EMC. Other utility libraries, such as produtil, are not part of NEMS.

The "caps" that translate native model component interfaces to NUOPC interfaces are not part of NEMS. They are logically housed in the repository with the model component, though this is not always the case.

Management Gaps

NEMS is designed to be a unified modeling system in which applications share components and infrastructure. NEMS applications consist of a set of SVN externals to specific revisions of model components and the NEMS infrastructure. There is no requirement that all applications use the same revisions of model components or NEMS itself.

NEMS is also a unified modeling system because of the underlying repository fabric that gives concrete meaning to the application/component matrix used to define NEMS requirements - see Appendix D. In this repository structure, each application that uses a component (including NEMS itself) uses it out of the same repository. This approach is designed to allow individual applications, which may have different components, coupling requirements, and timelines, to develop independently and at their own pace. This strategy has enabled diverse applications from different teams to be developed under NEMS.

However, a strategy is needed that defines when applications with changes on branches merge the changes to components and NEMS back to their respective trunks. This is necessary for coordination of development and is currently a major management gap.

A second management gap is that there is currently no standing decision-making body that spans NEMS applications and is able to resolve conflicts in requirements, assess design tradeoffs, and coordinate scientific and technical plans. Efforts to advance coupled systems would also benefit from partnerships that bring expertise in coupling science and techniques.

Unified Modeling Gaps

There are two main gaps in the NEMS software with respect to the unified modeling goal. The first gap has to do with restructuring the original NEMS software. When NEMS was first constructed, it included two atmosphere models, the Global Spectral Model (GSM) and the Nonhydrostatic Multiscale Model on the B-grid (NMMB), and a generalized atmosphere component that called each of these models. Bundling the general atmosphere component and atmosphere models with the coupling infrastructure created code dependency issues and complicated scripts.

During the last year, GSM and NMMB were removed from NEMS along with the generalized atmosphere component. Individual NUOPC caps were created for GSM and NMMB. This brought the treatment of the atmosphere in line with other components, such as ocean and sea ice. The gap is that not all applications have moved to the new atmosphere structure.

A second gap is that the changes made to the NEMS mediator for the regional system are still on a branch, and have not yet been merged back to the main mediator. The changes were not extensive.

The build system for NEMS (NEMSAppBuilder) is a single script that is convenient to use for standard cases. It is simply structured as a list of component build recipes that invoke the native component builds. Users have had trouble modifying it because changes require knowledge of the native component build options. It is currently in the process of being restructured and replaced by a team at EMC.

Mediator Gaps

It is possible to use multiple coupling approaches in NEMS. There are two NEMS mediators (couplers), the main NEMS mediator and a space weather mediator. The NEMS mediators are specializations of a NUOPC mediator class, and relatively small codes. Instead of a mediator, some simpler NEMS applications use only NUOPC connectors, which move data one way and perform basic functions like redistribution and grid remapping. These connectors can often be used as completely generic NUOPC library code.

The main NEMS mediator implementation has been used for UGCS-Seasonal, regional nest, and land/hydrology application milestones. The NEMS mediator follows a CESM-like approach, with all explicit coupling and no exchange grid.

If science leadership would like to move to a coupling approach that is scientifically and numerically more like GFDL's, with implicit coupling between some components and an exchange grid, or support both options, these options would need to be added to the mediator. The underlying ESMF/NUOPC infrastructure has an exchange grid class and can support implicit coupling. There may be implementation details that require development work. This could be considered more a science choice than a gap.

Coupling in 3D with the space weather mediator has been validated in the WAM-IPE space weather application in a one-way interaction. Work on two-way coupling has just started. This could be considered more of a development status than a gap.

Most of the applications will need to change the atmosphere model to the Finite Volume 3 (FV3) dynamical core. Again, this may be more of a development status than a gap.

Some of the applications have identified science improvements to the coupling that are described in the milestone documentation pages[10]. The most notable are an inconsistency in the flux calculation in UGCS-Seasonal, and a grid remapping issue in the regional nest code. A fix to the latter is being evaluated.

There are many feature additions and improvements that could be made to the NEMS mediators. None of these appear to be show-stopping in terms of NEMS functionality. Ongoing development or new requirements may expose additional gaps.

ESMF/NUOPC Gaps

There are advanced features of ESMF/NUOPC that make it well suited to NEMS. These are worth understanding as well as the gaps. The ESMF remapping approach, supported by an underlying 3D finite element mesh framework, allows for representing and remapping virtually any grids or meshes in parallel. The ability to change the run sequence of components at run-time, including slow and fast loops, also contributes to the framework's flexibility. The framework supports concurrent and sequential execution of components, including mediators.

A gap is that ESMF does not provide a 2nd order conservative interpolation scheme, which is preferred at GFDL. This is expected by the end of March 2017.

[10] Milestone documentation is under the Milestone Revisions header on the left navigation bar of the Coupled NEMS website: https://esgf.esrl.noaa.gov/projects/couplednems/

The ESMF and NUOPC team receives many feature requests and there are countless opportunities for improvement. None of these appear to be showstopping in terms of NEMS functionality, and the ESMF/NUOPC change review board periodically reviews and prioritizes these potential improvements. The framework software is actively being developed. Ongoing application development or new requirements may expose additional gaps.

Appendix F – Definitions

Additional definitions are in the *NEMS System Architecture Description* glossary, see:
https://esgf.esrl.noaa.gov/site_media/projects/nems-workshop/report_1610_system_architecture.docx

Application – a particular forecast target, characterized by the lead time, quantities to be predicted, and requirements for accuracy, acuity and reliability.

Component – “composable” software elements that have a clear function and interface - in coupled models, these are often a portion of the Earth system, e.g., atmosphere, ocean or land surface.

Cost effectiveness – the degree to which the costs in time, human or computational resources are in balance with the benefits (e.g. in improved forecast skill or advanced understanding of a particular development).

Ensemble – a collection of forecasts that are reasonably viewed as parts of a whole, e.g., by virtue of being equally probable by construction.

NCEP/EMC system architecture – a layered structure that is expected to encompass 1) a workflow environment that includes a user interface and database of previous runs and verifying analyses, 2) a prediction suite with a sequence of pre-processing, data assimilation, forecast, and post-processing components, 3) a model application layer with a coupling framework, a prescribed interface between atmospheric physics and dynamics, model components, and scripting, and 4) a layer of utilities and numerical libraries. This structure is shown in Figure 1.

Operational - functioning routinely with a fixed schedule and a well-defined set of deliverables, products or services .

Quality of architecture - the quality of a system architecture may have several attributes, including: adherence to software-engineering best practices; support for forecast skill; performance; clear, well-documented code, etc.

Software infrastructure - NCEP/EMC applications have been constructed using a software infrastructure that includes ESMF and NUOPC tools and standards. The software infrastructure is a set of technical building blocks that represents a wide range of implementation options. The software infrastructure should be distinguished from the system architecture - the latter defines what is built; the former is a set of tools for building it.

Stakeholders - the collection of interested individuals and institutions that have an interest in the outcome; for coupled prediction, stakeholders include either those who have the wherewithal to contribute to system advancement, have a mandate to produce forecasts in an operational manner or have requirements for the forecasts or forecast system.

System architecture –the fundamental organization of a **system**, embodied in its components, their relationships to each other and the environment, and the principles governing its design and evolution.

Appendix G – Selected recommendations from UCACN Model Advisory Committee

UMAC 2015 Points:

- Reduced complexity of the NCEP Production Suite.
- Rational, evidence-driven approach towards decision-making and modeling system development.
- A unified, collaborative strategy for model development across NOAA.
- NOAA needs to better leverage the capabilities of the external community.
- NOAA must develop a comprehensive and detailed vision document and strategic plan that maps out future development of national environmental prediction capabilities.
- Execute strategic and implementation plans based on stakeholder requirements.

UMAC 2016 Points:

- A useful and usable strategic plan is required that connects together NCEP, indeed NOAA modeling, from bottom to top
- Given the increasing complexity of forecasting systems expected in the next decade, NCEP requires an expertise base that is beyond its internal workforce, hence the importance of NCEP to more effectively work with the community: private sector, federal, and academic.
- Shared decision making is required, along with the tools necessary to support it
- Convection-allowing models and ensembles
- Salt- and fresh-water modeling capability

Appendix H – Preliminary Considerations for Data Assimilation Applications

As mentioned in Recommendation T4, the SAWG has begun discussion of the implications that data assimilation applications have for the system architecture. A very preliminary set of recommendations has been formulated as a basis for that discussion. Note that the recommendations enumerated below are not as yet a consensus view of the SAWG and will be refined for delivery in a later report.

Preliminary Recommendations for Data Assimilation Applications

	Importance	**Synthesis Recommendation - Data Assimilation Applications**
DA1	Essential	Ability to call forecast models (either individual or coupled domains) multiple times within the same executable, without significant overhead (i.e. without re-doing the setup).
DA2	Essential	Flexibility to call forecast models (either individual or coupled domains) with starting time, initial conditions, and forecast length controlled by the data assimilation.
DA3	Essential	Support a full data assimilation (DA) capability within a single DA system architecture, including DA in individual components, weakly coupled DA and strongly coupled DA.
DA4	Essential	Rely on requirements from the Joint Effort for Data assimilation Integration (JEDI) for DA system components.
DA5	Desirable	Provide tools and training to reach high-level understanding of DA concepts and needs across the modeling community.