

A Vision Paper for the Earth Prediction Innovation Center (EPIC)

Version 5.0 updated 5/28/19

Writing Team

Harry Cikanek¹, Patrick Burke², Derrick Snowden³, Ed Clark⁴, Brian Gross⁵, Ming Ji⁶, Ivanka Stajner⁵, Hendrik Tolman⁶, Curtis Alexander⁷, Stan Benjamin⁷, Jennifer Mahoney⁷, Whit Anderson⁸, Eugene Burger⁹, DaNa Carlis¹⁰, John Cortinas¹⁰, Sundararaman Gopalakrishnan¹¹, Tom Hamill¹², Jeff Whitaker¹², Robert Webb¹², Lou Wicker¹³, William Lapenta¹⁴

¹National Environmental Satellite, Data and Information Center for Satellite Applications and Research

²National Ocean Service, Center for Operational Oceanographic Products and Services

³National Ocean Service, Integrated Ocean Observing System

⁴National Weather Service, Office of Water Prediction, National Water Center

⁵National Weather Service, National Centers for Environmental Prediction, Environmental Modeling Center

⁶National Weather Service, Office of Science and Technology Integration

⁷Oceanic and Atmospheric Research, Earth System Research Laboratory, Global Systems Division

⁸Oceanic and Atmospheric Research, Geophysical Fluid Dynamics Laboratory

⁹Oceanic and Atmospheric Research, Pacific Marine Environmental Laboratory

¹⁰Oceanic and Atmospheric Research, Office of Weather and Air Quality

¹¹Oceanic and Atmospheric Research, Atlantic Oceanographic and Meteorological Laboratory

¹²Oceanic and Atmospheric Research, Earth System Research Laboratory, Physical Sciences Division

¹³Oceanic and Atmospheric Research, National Severe Storms laboratory

¹⁴Office of the Assistant Secretary of Commerce for Environmental Observation and Prediction

Background:

The Weather Research and Forecasting Innovation Act of 2017 (WRFIA), P.L. 115-25, instructs the National Oceanic and Atmospheric Administration (NOAA) to prioritize improving weather data, modeling, computing, forecasting and warnings for the protection of life and property and for the enhancement of the national economy. The National Integrated Drought Information System Reauthorization Act of 2018 (NIDISRA), P.L. 115-423, instructs NOAA to establish the Earth Prediction Innovation Center (EPIC) to accelerate community-developed scientific and technological enhancements into the operational applications for numerical weather prediction (NWP). Specifically, Section 4 of NIDISRA articulates Congress' vision for the EPIC program with the following responsibilities:

“Advancing weather modeling skill, reclaiming and maintaining international leadership in the area of numerical weather prediction, and improving the transition of research into operations by—

- A. leveraging the weather enterprise to provide expertise on removing barriers to improving numerical weather prediction;
- B. enabling scientists and engineers to effectively collaborate in areas important for improving operational global numerical weather prediction skill, including model development, data assimilation techniques, systems architecture integration, and computational efficiencies;

- C. strengthening the National Oceanic and Atmospheric Administration’s ability to undertake research projects in pursuit of substantial advancements in weather forecast skill;
- D. utilizing and leverage [sic] existing resources across the National Oceanic and Atmospheric Administration enterprise; and
- E. creating a community global weather research modeling system that—
 - i. is accessible by the public;
 - ii. meets basic end-user requirements for running on public computers and networks located outside of secure National Oceanic and Atmospheric Administration information and technology systems; and
 - iii. utilizes, whenever appropriate and cost-effective, innovative strategies and methods, including cloud-based computing capabilities, for hosting and management of part or all of the system described in this subsection”.

The amended WRFIA language for EPIC describes NWP in the global scale context. Therefore, the near-term program formulation will be focused on the end-to-end operational Global Forecast System (GFS) composed of the Global Data Assimilation System (GDAS), the model and Unified Post Processor (UPP). It is expected that the scope of EPIC will expand to include other operational model applications and mission priorities outlined in the 2017 WRFIA such as convective allowing models (i.e., High Resolution Rapid Refresh) and fully coupled seasonal-to-subseasonal (S2S) forecast systems (i.e., the Climate Forecast System, National Water Model, Ocean Forecast Systems). NOAA also recognizes the need to extend R2O2R standards and practices to all disciplines within the agency and to create the advanced capability needed to extend weather model output downstream to water, land, living resources, and human communities.

The transition of research into operations and operations into research (R2O2R) has been studied and reported on for decades. NOAA established testbeds and proving grounds to facilitate the orderly transition of research capabilities to operational implementation through development testing in testbeds, and pre-deployment testing and operational readiness/suitability evaluation in operational proving grounds ([NOAA Testbeds and Proving Grounds Portal](#)). At least two of the testbeds are focused on NWP and based on interagency agreements between NOAA, the National Aeronautics and Space Administration (NASA), the United States Air Force (USAF), The United States Navy (USN) and the National Center for Atmospheric Research (NCAR) to specifically address: 1) satellite data assimilation, via the Joint Center for Satellite Data Assimilation (JCSDA), established in 2001; and 2) mesoscale modeling, via the Developmental Testbed Center (DTC), established in 1999. The amended WRFIA states EPIC will utilize and leverage existing resources across NOAA. Given the initial scope of EPIC is focused on global NWP it must effectively leverage activities under the JCSDA and DTC to be successful.

A fundamental question in the near term, is how EPIC will accelerate the rate of transitioning innovative research and development into NOAA NWP operations. A related question is how EPIC will build upon NWP success to transition research into other NOAA operational numerical guidance applications (e.g. climate, hydrologic, ocean). There are several important cultural, organizational and technological developments that have been established in the past

several years within and external to NOAA that position EPIC to be successful in the near term and in the long term:

- 1) The National Science Foundation (NSF), NOAA, NASA, USN and USAF increased interoperability and shared use of component models through co-development of common model architectures and infrastructure.
- 2) The National Centers for Environmental Prediction (NCEP) is simplifying the National Weather Service (NWS) operational production suite based on the recommendation of the University Corporation for Atmospheric Research Community Advisory Committee for NCEP (UCACN) Model Advisory Committee (UMAC).
- 3) NOAA adopted a unified modeling strategy and established the Unified Forecast System (UFS) in partnership with the community which builds off of the common model architecture and infrastructure co-developed with other Federal agencies.
- 4) NCAR and NOAA recently signed a Memorandum of Agreement (MOA) to continue co-development of the common model infrastructure to support the UFS.
- 5) NOAA developed a Strategic Implementation Plan (SIP) working with federal partners, academia and private industry to define scientific and development priorities for the operational applications of the Unified Forecast System (UFS) through 14 working groups.
- 6) The NOAA Oceanic and Atmospheric Research (OAR) Office of Weather and Air Quality (OWAQ) and the NWS Office of Science and Technology Integration (OSTI) have strengthened coordination of funding programs across organizational boundaries to address scientific priorities identified in the SIP working groups through coordinated announcements of opportunity.
- 7) NOAA Fisheries published a roadmap for implementing Ecosystem Based Fisheries Management, a Climate Science Strategy, and a Fisheries Stock Assessment Improvement Plan. All of these can benefit from EPIC's approach and from associated advancements in coupled end-to-end models, particularly at S2S scales.
- 8) The Alaska Climate Integrated Modeling (ACLIM) involves multi-disciplinary NOAA and CI scientists demonstrated the merits of a unified approach to modeling the earth system from regional circulation to ecosystems and the economics of fisheries.
- 9) Cloud-based high-performance computing environments have matured to the point to enable efficient and effective processing of end-to-end NWP systems.
- 10) Within the National Water Initiative, NOAA has recognized the value of coupling NOS expertise in service delivery within a community modeling framework with NWS expertise in operational modeling and forecasting. This combination has allowed us to build towards a best in class, science based service delivery framework.

Defining the EPIC Vision:

A workshop was held on 6-7 March to begin the EPIC project formulation. Twenty NOAA Federal employees participated, representing over 400 combined years of experience in numerical modeling and/or the transition of research into operations (R2O). The purpose of the workshop was to define the scope of EPIC and explore options of how it will enhance model development and transition processes within the agency. Participants agreed EPIC must provide a framework to accelerate R2O through community-developed enhancements into the UFS and its operational applications. The UFS is a community-based, coupled comprehensive Earth

system model-based analysis and prediction system designed to meet NOAA's operational forecast mission to protect life and property and improve economic growth.

Seven 7 core investment areas for EPIC were identified:

1. Software engineering
2. Software infrastructure
3. User support services
4. Cloud-based high performance computing
5. Scientific innovation
6. Management and planning
7. External engagement and community

The workshop participants agreed the UFS must be scientifically credible, documented, supported, computationally flexible, managed and available to the entire scientific community for use in fundamental research, development of system improvements and operational applications. This will require investments in software engineering, software infrastructure, user support services and the ability to run the system on cloud-based high performance computing. It was also agreed EPIC must provide a management function to communicate, adhere to and guide the [R2O2R transition process](#). Finally, EPIC must effectively establish clear lines of communication and coordination between scientists within and external to NOAA to deliver scientific innovation within the UFS framework.

The purpose of this Vision Paper is to: a) validate the EPIC concept defined herein with NOAA leadership; b) provide the foundation and guidance for the EPIC Program to begin implementation; c) socialize the concept with the prospective partners and modeling community, and; d) communicate with Administration and Congressional stakeholders to respond to EPIC authorizing language.

The following sections provide descriptions of the major components within the priority investment areas as identified at the workshop. The last section provides a summary and near-term roadmap for next steps.

1. Software and Performance Engineering

Modeling systems for NWP and other components of the full-earth system within the UFS, and the operational applications derived from them, are becoming ever more complex. The days when a single scientist could understand the entire system are over. In order for innovative research to be integrated and tested in the UFS quickly, algorithms must be encapsulated into sections of code that have well-defined interfaces. This way, a scientist can focus on his or her particular area of interest without having to know the details of how other parts of the model function. In the computer science industry, this design principle is known as *separation of concerns* (SoC). A key part of EPIC will be to bring modern software engineering practices, such as SoC, to bear on the model development process. EPIC will employ professional software engineers to work directly with scientists in an integrated cross-disciplinary environment, to ensure that the UFS code is easily maintainable, extensible and performance optimized. A key part of this will involve ensuring that all code modules include tests and can

be reused across components of the system wherever possible. In addition to working directly with scientists to integrate new algorithms, EPIC software performance engineers will be needed to continually refactor, modernize and optimize the existing code as software and computing platforms evolve.

Another key role for the EPIC software engineering staff will be to ‘containerize’ the UFS software. Containers provide a packaging mechanism in which applications can be abstracted from the environment in which they run, whether that is a scientist’s laptop, a university HPC system or in the cloud. Containerization is a key aspect of the SoC design principle, since it allows developers to focus on their scientific algorithms, separate from platform-specific compatibility and dependency issues.

Testing is critical to ensure the software quality for newly modified code, assure robustness, and verify performance. EPIC’s testing infrastructure should encompass regression, hierarchical unit, quality assurance, code coverage and performance testing. Unit tests will be required for any new code to be accepted into the UFS. Continuous Integration (CI) approaches that *automatically* build and test all proposed code changes will be an essential component of the EPIC testing infrastructure.

Talented software engineers are in high demand in the tech industry, and salaries reflect this. In order to recruit the necessary talent, EPIC will need to devote a significant portion of its budget to software engineering, more than has traditionally been done within NOAA. To stimulate knowledge infusion and foster effective collaboration, there should also be cross-training opportunities, where software engineers are embedded with the scientists to understand the science, and share modern software engineering best-practices. This will foster better communication while promoting greater staff investment in the core goals of EPIC.

2. Software Infrastructure

Software infrastructure is defined as the foundational code and processes upon which applications are built and managed. The UFS is an end-to-end system for earth system prediction that encompasses the assimilation of observations (i.e., efforts such as the Joint Effort for Data assimilation Integration (JEDI)), development of initial conditions, execution of the prediction model, post-processing and analysis of model output, verification of results, and the delivery of products. Software infrastructure must support this entire process.

The Memorandum of Agreement between [NCAR and NOAA for the co-development of common model infrastructure](#) outlines areas of model development that help define the scope of software infrastructure considered here. These include:

2.1 Open, managed, authoritative repositories

Community involvement will play a critical role in the development and the progression of software from concept through acceptance into operations. This requires not only a robust repository infrastructure, but also a dedicated staff to manage version control. An open, authoritative and well-managed repository infrastructure, based on Git and the associated cloud-based GitHub. Git is a distributed version-control system for tracking changes in source code during software development. It is designed for coordinating work among programmers, but it

can be used to track changes in any set of files. Repositories will exist for umbrella applications and different components. Each component repository will contain the source code for a unique component of the UFS application. The umbrella repository will contain the policies, documentation, configurations, and testing protocol required to link to the individual component repositories, which, when brought together, define a given UFS application. The UFS Repository Management Plan provides one example of such a repository governance.

2.2 Workflow Management

In the context of this document, workflow refers to all infrastructure, code, and datasets needed to configure, build, test, and run an end-to-end forecast system. This common workflow infrastructure should be developed and maintained to facilitate community research and accelerate the adoption of new research into operations. Examples of current workflows include CIME from NCAR, Rocoto from EMC, and JEDI from the JCSDA. The intent is to make the workflow portable enough that it successfully executes on a few commonly used computing platforms

2.3 Intra- and Inter-Component Coupling

The Earth System Modeling Framework (ESMF) architecture and the National Unified Operational Prediction Capability (NUOPC) layer are well-established community software packages. Building on the NUOPC infrastructure, the NCAR- and NOAA-developed Community Mediator for Earth Prediction Systems (CMEPS) is a new framework for inter-component coupling of community models that can ultimately accommodate multiple model coupling strategies. EPIC should develop, support and maintain a common CMEPS mediator.

3. User Support Services:

One of the key areas for EPIC success is establishing comprehensive user support services to enable a broad and vibrant community of UFS users and developers. User Support Services (USS) will nurture a collaborative, community-based framework to support, advise and educate a cadre of system and software users as well as developers. A valuable use case for USS is the [Graduate Student Test](#), which demonstrates that the code is well structured and usable, such that, in a matter of hours, a graduate student can install and test the code, then begin running the code for development purposes.

3.1 Technology

Initial USS is a near-term (0-3 years) high priority for EPIC. As EPIC matures, these functions are envisioned to extend in depth, scope, and expertise. The following core technology capabilities will be made publicly available:

- Code portability, as described in Section 1, to include the workflow for running the code on computing systems accessible to the user community.
- The code in a managed public authoritative repository, as described in Section 2.
- Test-cases, with supporting data, to ensure that the system compiles and runs correctly.
- A verification package for the system, with regression tests and a set of standard metrics.

Thus, the core deliverable in the first three years is an end-to-end UFS software system for the user communities that is portable (platform-independent), user-friendly, and supported by EPIC resources and staff. Components of this platform-agnostic system are currently being developed

using Docker containers for community modeling systems. Lessons learned from investigative NOAA computing efforts, such as the Big Data Project and ongoing collaborations with Azure and AWS cloud services, will also be applied to inform the development of these capabilities. This end-to-end system will be a software capability that includes the acquisition of observations, the data assimilation system, the model forecast system, the post-processing, and a verification capability.

3.2 Documentation

Comprehensive code and science-based documentation is essential for UFS usability. EPIC will deliver a comprehensive set of documentation (both ‘in-line’ to the code and external) that supports scientific familiarization and understanding of model code by researchers and developers. The documentation will be maintained dynamically by both the user community and EPIC support staff to ensure it stays up-to-date as the code is updated or revised. EPIC will enable the community to grow (and act as a force multiplier), teach themselves, and foster a more collaborative approach to code development. The documentation will address all four core technology capabilities, including: how the code is managed and code versions are controlled; the types of administrative code management activities that are performed; best practices that will be enforced to ensure model code portability to a variety of computing platforms; and the test cases and verification packages used. This initial large investment in comprehensive documentation likely will “pay” for itself in the long-term reduction of effort needed to support the UFS system through staffed help desks and moderating and responding to user forums. To help introduce users to the system, along with the support staff and documentation, regularly scheduled user tutorials and workshops (both in-person and online) will be provided by EPIC, particularly during the first years of EPIC.

3.3 Community Development Support

To facilitate the initial growth of the user community, it may be beneficial to establish a user help desk. EPIC USS will promote collaborative user interfaces, such as virtual forums, and adopt best practices from the software development and science communities. Such technological interfaces must be compliant with Department of Commerce (DOC) and NOAA policy, and may include:

- Documentation functions within GIT repositories, as recommended in the Software Infrastructure plan.
- Real-time community dialog, such as a Slack channel.
- Collaboratively developed, modified and dynamically maintained web-based documentation (i.e. wiki).
- Online tutorials, e.g., virtual Hang-outs or YouTube (educational) videos.

EPIC will establish routine in-person user-focused workshops and sessions or forums at meetings, such as AMS and AGU conferences. As the community grows, EPIC should consider developing Massive Online Open Courses (MOOCs). USS should grow the community and foster partnerships and collaborative activities with the following:

- Cooperative Institutes and NCAR.
- NOAA Line Offices, other Federal agencies and private industry.
- Academic institutions and possible including UFS in graduate curricula.

4. Cloud-Based High Performance Computing:

The EPIC high-performance computing (HPC) strategy has two core components designed to promote rapid advancement of the UFS system by NOAA scientists and users from outside NOAA. The first component will enable a scientist to rapidly run models and conduct research employing the current version of the operational NWP system run at NCEP and, potentially, other operational centers, establishing the *first* operational end-to-end NWP system available to weather researchers in the United States. The second component lowers the compute-barrier often preventing researchers from using NWP systems by providing a software system that runs on a variety of hardware platforms, including the cloud, university HPC, as well as traditional HPC available within NOAA, Department of Energy (DOE), NASA, and NCAR organizations. This capability will be crucial for rapid adoption of the system.

4.1 EPIC and a UFS Compute Strategy

The chief barrier to broadening the base of earth system modeling research is the lack of research HPC resources within organizations (e.g., NOAA). For example, global and regional high-resolution NWP is a resource intensive endeavor that stresses most computing infrastructures (e.g., central processing units (CPU), input/output (I/O), and disk capacity). Traditional NWP research has occurred only on “big iron”, i.e., a large system internally managed by organizations who periodically can invest large resources into their HPC. HPC availability is a limited resource, which greatly limits the number of users who can do NWP research. For government labs, another major barrier is access, due to the security clearances now required for accessing their internal and agency-wide shared systems. Very few short-term (less than 1-year) projects from non-Federal employees (e.g., university faculty and students or private industry personnel) are ever realized. Researchers with good ideas cannot get onto a machine and test their ideas before having to move onto their next project.

The success of the UFS, therefore, may ultimately depend on its own “user scalability”. We believe that scaling up UFS research activities requires the UFS to be portable across multiple computational architectures to enable access for a much-larger user base. This approach is used by NCAR for the Weather Research and Forecast (WRF) model. The WRF runs on a large number of platforms, ranging from the laptop to high-end HPC at the national centers. Importantly, the rapid growth and proliferation of cloud computing resources over the past decade now provides a new way forward that an end-to-end UFS should leverage. EPIC should focus on providing a UFS that can be easily run on cloud computing HPC systems that are currently most suited for large-scale earth system prediction. NOAA should also invest in software engineers (sec. 1) to work with NWP scientists, to develop a portable system that runs with ease on nearly any platform having sufficient resources. Over the next five years, NOAA, through the Office of the Chief Information Officer needs to fully develop and implement a cloud computing strategy for research and development to properly plan for and develop the capability. However, this does not preclude the need to maintain our current internal R&D and operational HPC “big-iron” - these computing resources will always be needed for high-readiness-level (RL) projects transitioning into NCEP operations. Maintaining a balanced portfolio, with growth in both cloud computing and NOAA’s internal HPC system, will help

mitigate risks to our computing needs, given that cloud-computing vendors are very much market driven.

4.2 Strategic Partnerships and a Next-generation Earth Prediction Capability

The development of a next-generation prediction system based on exascale computing (i.e., computing systems capable of at least one exaFLOPS, or a billion billion calculations per second) is a grand challenge activity that several agencies (DOE, NASA, NOAA) have begun researching. There exist many challenges in both hardware and software to move beyond the current performance levels (i.e., $\sim O(10)$ Petaflops) of existing UFS applications. Exascale computing is initially likely to represent an aggregated compute capability for large high-resolution ensembles or coupled systems. Whatever the path forward, EPIC will need to develop strong partnerships with other agencies to develop effective computing strategies to reach the exascale performance for comprehensive earth system prediction capabilities that will be needed to assess and forecast weather and climate in the latter half of the 21st century. These activities will occur higher up the R2O funnel than traditional NOAA research initiatives, requiring extending our thinking about the UFS to a multi-decade activity.

5. Scientific Innovation:

The “I” in EPIC stands for innovation. Therefore, EPIC must be designed to stimulate innovation and the transition of research into UFS operational applications. Workshop participants generally agreed on three points. First, NOAA research-to-operations programs (e.g., JTTI and NGGPS) are most concentrated on high RLs, i.e., finalizing relatively mature applications for operational use. Second, there need to be periodic research and development (R&D) portfolio evaluations that include assessments of low-RL project viability, maturity, and priority for advancing to the next development stage, along with a formalized, evidence-based “stage-gate” process in EPIC and throughout NOAA. Third, what NOAA funding there is for low RL activities originate in other agencies or in NOAA laboratories. There are no formal mechanisms in place for sustaining these promising projects as they progress to higher RLs. The formation of EPIC is a chance to reconsider how R&D across RLs can be managed more effectively to maintain a stream of innovations into NOAA’s operational mission. There is an opportunity to improve this process and define a more specific role for EPIC (consistent with NOAA R&D policies - see [here](#) and [here](#)) with the update of the NOAA R&D plan for 2020-2026.

The precise role for EPIC in the scientific innovation process is still up for discussion. The figure below illustrates several possible non-exclusive options within the paradigm of the NOAA R2O funnel; however, more ideas are welcome. All options are based on the premise that EPIC will deliver a robust community prediction system based on the UFS and that enterprise partners will be using the UFS for their low-RL R&D. In panel (a), using evidence-based decision-making procedures and evaluators that span the R&D funnel, EPIC would foster the improvements coming from the broader weather enterprise and recommend what is ready for further development at mid RLs in, for example, OAR, NESDIS, or NOS for potential transfer to operations. In panel (b), EPIC provides a similar functionality, but evaluates contributions from NOAA and the wider enterprise at mid-to-high RLs, deciding what system improvements are ready for pre-parallel testing. In panel (c), EPIC would provide the scientific direction to

addressing scientific grand challenges, shepherding promising projects from their infancy through to operational implementation.

In order for EPIC to stimulate and increase the rate of innovation into NOAA operations, it must deliver the community-based UFS, connecting high- and low-RL projects within the funnel. To eliminate discontinuities (i.e., hand offs) between research, development, and operational elements, EPIC must also provide a communication continuum across the RL's and organizations within the funnel. It is envisioned that EPIC should provide a collaborative environment capability, where scientists from NWS (i.e., Environmental Modeling Center, Meteorological Development Laboratory), NESDIS, OAR, NOS, and NMFS labs, academia and private industry can gather to conduct innovative co-development and associated testing. The collaboration capability will be designed to stimulate innovation, allow scientists to conduct higher-risk work, and increase NOAA's ability to engage world-class scientists and software engineers.

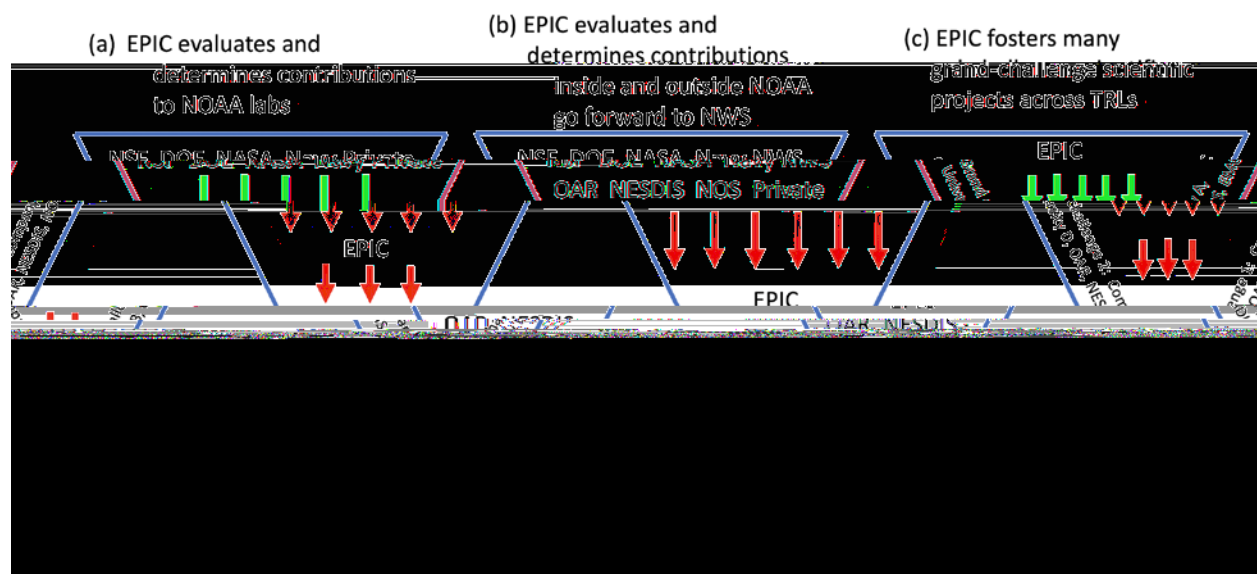


Figure 1: An illustration of three potential roles for EPIC in fostering scientific innovation in the development of the UFS. See associated text for related discussion.

6. Management and Planning:

The EPIC “Management and Planning” investment area provides the leadership and administration for the two functions central to EPIC success. The first function is to build and sustain EPIC as a 21st Century distributed community-model development capability by providing leadership and support to efficiently and effectively execute the other 6 components and functions that comprise EPIC. To facilitate this, NOAA has placed the EPIC Program within OWAQ.

During the early phase of formulating and defining EPIC, options for acquiring and managing the components will be assessed. A key governance element of this will be to establish a management structure that enables a NOAA R2O2R framework for earth system prediction models that effectively harnesses the community advances in understanding and modeling while

providing a value proposition to attract partners. This capability will be resourced by integrating NOAA research labs and programs (OAR, NWS, NWS, NOS), developmental organizations (e.g., EMC, NOS/OCS, and NESDIS/STAR), and Operational Line Offices (NWS, NOS, NMFS, NESDIS) within EPIC. NOAA-wide commitment to and support of matrix-managing the EPIC-leveraged pool of resources and capabilities is necessary. This approach will ensure that EPIC priorities will be integrated into the much-larger NOAA portfolio of activities that are critical to achieving EPIC goals.

The second function will be services required to establish a stakeholder engagement process, workshops, training and forums to identify, articulate and sustain NOAA operational needs, plans and roadmaps. The foci underlying these services are to: a) drive strategy, priorities and funding for model research and development programs and researchers and developers; b) communicate grand challenges in prediction to the research and development communities; c) implement internal evidence-based decision-making processes for evaluating potential system improvements; and d) guide associated HPC planning. As a sustained partnership, NOAA will need to plan for investing in the infrastructure required to support easy internal and external community access and use of model code, software engineering at all development stages, and high-performance computing capabilities enabling broad community participation. Within the strategic management component, EPIC provides NOAA the facility to formulate and communicate operational needs over short to long term to the developer/researcher community, while also providing the facility for the developer community to formulate and communicate roadmaps for the art of the possible over time, enabling strategic planning and resource allocations for research and model development/improvement programs and future operational model suite improvements.

The initial focus for EPIC implementation will be on NWP and the UFS. As the UFS evolves toward full coupled Earth System Modeling, providing guidance for operational applications to support the coastal marine environment, hydrological forecasting overland and at the coasts, and the sustainable exploitation of living resources will be essential to the overall success of EPIC. Evolving EPIC to fulfill this potential will be a responsibility of the management and strategic planning function and pursuing these strategic enhancements will require input from across NOAA.

7. External Engagement and Community:

The vision for EPIC presumes that model development through a community of NOAA internal and external partners will yield an increase in the rate of improvement in its operational systems. Successfully attracting external partners to participate in a community development effort requires an engagement strategy focused on identifying willing partners, engaging them to understand their needs, and incentivizing them by providing an environment that will enable fulfilling those needs while also accelerating development of earth system models NOAA depends on for operational service delivery.

Delivering a new and fresh approach to engaging the external community requires a different thought process for conducting business. Over the past few years, NOAA has set out to develop an integrated external engagement approach that allows partners to be deeply involved in decision making across NOAA's modeling infrastructure (i.e. Unified Forecast System Steering

Committee). EPIC will build upon these efforts in shared decision making and priority setting by creating an environment where community members can actually work on the same problems using the same tools (predominantly model code). This environment will allow the development community to easily share knowledge, tools, software, training, and computational architecture in the form of commercial cloud computing. EPIC is, in part, a program to organize and enable community collaboration, along with a computational system that can serve as infrastructure for this community. For EPIC to be successful, the community must, at all stages of development, have a clear voice in the design and execution of the program and the computational system.

Community engagement activities can be grouped in support of two main goals. First, engagement is a critical piece of a system development life cycle, serving to identify and understand the requirements for a functioning system. Feedback from potential stakeholders helps inform the design of the program and the computational system, ensuring both meet research and operational needs. Typically, successful software development projects consistently conduct this type of engagement or customer interaction throughout the project lifecycle, not just at the beginning.

The second element of engagement aims to ensure that the developer community (scientists, software engineers, etc.) remains vibrant and active. This activity fosters community by creating opportunities for members to interact and share ideas, knowledge, and tools. The scoping committee agreed that EPIC should support several activities aimed at building and maintaining the developer community, including:

- ensuring a reliable and regularly updated web presence,
- instituting a visiting scientist program,
- supporting hackathons and code sprints to galvanize developers around a problem,
- embedding software engineers with scientists to encourage cross-disciplinary input into model development, and
- encouraging best practices for software engineering.

NOAA must avoid a “build it and they will come” mentality. Rather, from inception, we should ingrain community engagement into the core principles of the EPIC program and accept responsibility for maintaining active and persistent efforts to engage the community on their terms, make the case for exploiting EPIC to enable community goals, and be open to feedback that might change EPIC’s evolution. Cultivating a culture of community engagement is well worth the effort, given the potential for improvements to our operational forecast portfolio.

8. Summary and Next Steps:

The reauthorization of the Weather Research and Forecasting Innovation Act in January 2019, P.L. 115-423, instructs NOAA to establish the Earth Prediction Innovation Center (EPIC) to accelerate community-developed scientific and technological enhancements into the operational applications for numerical weather prediction (NWP). The initial focus is on the operational global medium range NWP application known as the Global Forecast System (GFS).

A workshop was held on 6-7 March to begin the EPIC project formulation within NOAA. Twenty NOAA Federal employees with more than 400 combined years of experience in numerical modeling and/or the transition of research into operations participated. Workshop

participants agreed EPIC must provide a framework to accelerate R2O through community-developed enhancements into the UFS and its operational applications. Seven priority investment areas were identified in order for EPIC to be successful. Software Engineering and Software Infrastructure were identified as the top two investment areas. This finding is consistent with the work identified within the Memorandum of Agreement between [NCAR and NOAA for the co-development of common model infrastructure](#).

The OAR Office of Weather and Air Quality (OWAQ) manages the U.S. Weather Research Program (USWRP), and is responsible for executing the EPIC program. In FY19, the primary activities will be managed by the OAR Program Manager for the Next Generation Global Prediction System (NGGPS), and the OWAQ Director will have ultimate oversight of the program. For FY19, the EPIC program will leverage existing NWS NGGPS management support to assist with EPIC program management. Funding for FY19 activities associated with EPIC activities will be supported from several sources: (1) FY18 Disaster Relief Supplemental funding; (2) the Joint Technology Transfer Initiative (JTII); (3) USWRP, and (4) other base resources within OAR and NWS.

The EPIC program is currently developing a strategy to effectively and fairly engage the external NOAA community. It is envisioned NOAA will issue a Request for Information (RFI) in the June 2019 timeframe to solicit input from the community on how best to meet the goals of EPIC. The RFI will be analyzed to inform a community engagement workshop to be held in early August 2019. The exact structure of the workshop depends on the input received via the RFI. The input received via the community workshop may be used to formulate a Request for Proposals (RFP) to be released Q1FY20 with awards implemented in Q2FY20.