Triennial Report on Computing Resources Prioritization

Developed pursuant to: Section 108 of the Weather Research and Forecasting Innovation Act of 2017, (Public Law 115-25, as amended by Public Law 115-423)

Zachary G. Goldstein, Chief Information Officer and Director
High Performance Computing and Communications
National Oceanic and Atmospheric Administration

Neil A. Jacobs, Ph.D.,
Assistant Secretary of Commerce for Environmental Observation and Prediction,
Performing the Duties of Under Secretary of Commerce for Oceans and Atmosphere

TITLE I, SECTION 108 OF THE WEATHER RESEARCH AND FORECASTING INNOVATION ACT OF 2017 (PUBLIC LAW 115-25) INCLUDED THE FOLLOWING LANGUAGE

"Not later than 1 year after the date of enactment of the National Integrated Drought Information System Reauthorization Act of 2018, and triennially thereafter until the date that is 6 years after the date on which the first report is submitted, the Under Secretary, acting through the Chief Information Officer of the National Oceanic and Atmospheric Administration and in coordination with the Assistant Administrator for Oceanic and Atmospheric Research and the Director of the National Weather Service, shall produce and make publicly available a report that explains how the Under Secretary intends——

> *(1) to continually support upgrades to pursue the fastest, most powerful, and cost-effective high performance computing technologies in support of its weather prediction mission;*
> *(2) to ensure a balance between the research to operations requirements to develop the next generation of regional and global models as well as highly reliable operational models;*
> *(3) to take advantage of advanced development concepts to, as appropriate, make next generation weather prediction models available in beta-test mode to operational forecasters, the United States weather industry, and partners in academic and Government research;*
> *(4) to use existing computing resources to improve advanced research and operational weather prediction;*
> *(5) to utilize non-Federal contracts to obtain the necessary expertise for advanced weather computing, if appropriate;*
> *(6) to utilize cloud computing; and*
> *(7) to create a long-term strategy to transition the programming language of weather model code to current and broadly-used coding language."*

THIS REPORT RESPONDS TO THE APPROPRIATE REQUEST.

TABLE OF CONTENTS

I. EXECUTIVE SUMMARY

Maintaining and growing High Performance Computing (HPC) capability is one of the National Oceanic and Atmospheric Administration's (NOAA) highest priorities. Advancements in operational forecasts, predictions, and projections cannot occur without sustained investment in NOAA's HPC resources. NOAA's HPC resources are critical to operational forecasts, predictions, projections, and applications, and to the research and development (R&D) modeling efforts to improve those operational products. Meeting NOAA's mission requires a sustained growth and balance of operational HPC and R&D HPC. Both systems are required to maintain an effective and evolving program that is capable of delivering model-based products important to the Nation today and into the future.

II. INTRODUCTION

This report responds to Title I, Section 108 of the Weather Research and Forecasting Innovation Act of 2017 (Public Law 115-25, April 18, 2017, hereafter referred to as the "Weather Act"), as amended by the National Integrated Drought Information System Reauthorization Act of 2018 (Public Law 115-143, Section 5(b), Jan 7, 2019, 15 U.S.C. § 8518).

III. NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION – HIGH PERFORMANCE COMPUTING OVERVIEW

The advances in scientific understanding, applications, and predictions achieved by NOAA for its mission goals are critically dependent on sustained access and improvements to HPC. New applications are continually under development to address the pressing scientific challenges that can accelerate NOAA towards achieving its main mission of protecting life and property underpinned by credible, innovative science. Weather and water prediction based on numerical models supports an increasingly weather-vulnerable economy and society. Climate simulations, predictions, and projections inform science, service, and stewardship decisions that assess the current and future states of the climate system and potential impacts. The R&D portion of NOAA's HPC program provides increasingly powerful applications, as well as serves as a technology infusion to maintain cutting-edge relevance, providing advances across all of NOAA's scientific pursuits.

Through a combination of observations, increasing modeling detail and efficiency, along with the computational capability to ingest ever-increasing amounts of data, the R&D HPC program provides a path toward increased understanding of the environmental process and interactions. This better understanding provides a roadmap of where to target resources for future development of new technologies and techniques that can enhance the R&D HPC program, as well as improve downstream modeling and operational predictions. Improvements in this value chain are not limited to computational improvements, but are also applied to a nationwide system with many input/output points to increase efficiencies and capabilities. Enhanced network speeds, storage capabilities, data processing, end-user engagement, and product delivery are constantly evaluated, adjusted, and/or improved to maintain a world-class system consistent with current industry practices to efficiently provide the best science-based service output.

NOAA's work, across the agency, is often computationally extensive, expensive, and demanding because it must properly account for the complex interactions of the atmosphere, ocean, land surface, cryosphere, chemically active atmospheric constituents, biogeochemical cycles on land and in the ocean, and terrestrial and oceanic ecosystems. Furthermore, the applications must address accurately the information needs from local-to-global spatial scales, as well as timescales from seconds to century.

NOAA anticipates that its mission requirements will expand even more significantly than over the past decade, requiring additional high-caliber HPC capacity and use of cloud computing resources. Higher resolution and more complex models that capture the realism of the Earth System, combined with the use of ensembles (a set of computational algorithm configurations to characterize and quantify uncertainties), will require significantly enhanced HPC capabilities. These capabilities will also require new cutting-edge approaches in software, hardware, and data assimilation, management, transmission, and storage.

NOAA manages and operates major investments for operational and R&D HPC. These investments are managed as an integrated enterprise by the NOAA Office of the Chief Information Officer's (OCIO) High Performance Computing and Communications Program (HPCC). NOAA's HPC Board Allocation Committee (composed of representatives from all of the modeling centers across the agency) provides strategic guidance and oversight for the execution of this enterprise and the allocation of HPC resources. With the oversight of OCIO, integrated teams within NOAA's National Weather Service (NWS) and Office of Oceanic and Atmospheric Research (OAR) manage the day-to-day operation of these systems.

The NOAA Weather and Climate Operational Supercomputing System (WCOSS) provides reliable operational HPC capabilities essential to run real-time numerical models generating millions of weather guidance product outputs daily. This guidance is incorporated into public and private sector forecast processes to protect our Nation's lives and livelihoods. Since this investment is critical to the Nation, NOAA operates geographically dispersed identical primary and backup systems in Reston, Virginia, and Orlando, Florida.

The NOAA R&D HPC system enables significant improvements in weather, water, ocean, and climate research and modeling, as well as the increasing performance of the HPC systems themselves. NOAA operates R&D computing centers located in Fairmont, West Virginia, and Boulder, Colorado, and works with a Cooperative Institute-based computing center in Starkville, Mississippi, primarily driving the development of weather and seasonal to inter-annual models that will transition to operations. NOAA's research computing operated by Oak Ridge National Laboratory primarily supports improvements in the skill, resolution, and complexity of models used for "seamless" Earth System research, understanding, predictions, and projections – from daily weather to subseasonal to seasonal to decadal timescales and from county to regional to global space scales. This includes extreme events and transitions to operations and applications.

All of the systems – both operational and R&D – make up a balanced HPC portfolio that addresses the needs of a variety of scientific modeling applications. Each system is tailored to specific scientific, service-based workload requirements, ensuring that it is cost-effective. Operational applications, such as the operational weather models run by NWS, require highly reliable and available systems. These operational systems tend to cost more on a per-unit basis than R&D systems because of the stringent high availability requirements. By the same token, NOAA's R&D workloads have differing characteristics, allowing for greater flexibility in the use of systems with less availability and emerging technologies. For these reasons, NOAA's R&D HPC systems tend to cost less than similarly sized operational systems.

IV. NOAA'S HIGH PERFORMANCE COMPUTING UPGRADE STRATEGY

In Fiscal Year (FY) 2020, NOAA released a set of six strategies to accelerate the implementation of the most effective science and technology applications to advance NOAA's mission. These strategies are Cloud, Artificial Intelligence (or AI), 'Omics, Unmanned Systems, Citizen Science, and Data. NOAA's HPCC program supports NOAA's Cloud Vision of an accessible, scalable, secure, flexible, agile, cost-effective, efficient, highly reliable, and integrated suite of cloud computing products and services to fulfill its mission. Cloud-based HPC services will provide a flexible environment to extend NOAA's current capacity and accommodate emerging needs, serving as an intelligent platform to increase efficiency of the entire HPC resource. In FY 2020, NOAA has invested over $25 million into expanding cloud computing resources, with over $7 million of this directly supporting modeling and R&D HPC use cases on cloud resources.

NOAA's strategy to continually upgrade its HPC resources is generally the same across operational HPC and R&D HPC. NOAA upgrades HPC technology as resources allow and as a way to manage its ongoing compute operations and maintenance costs. As HPC hardware ages, the maintenance costs increase substantially as warranties and maintenance contracts expire. Over the past decade, NOAA has adjusted this strategy in response to technological advancements delivered by industry and to periodic influxes of funding from supplemental appropriations (e.g., disaster funding from hurricane supplementals).

NOAA has traditionally been able to replace its HPC systems every 3 years under a generally level budget because hardware technology advancements enabled a doubling of application performance every 2-3 years for a similar price-point without major re-work to NOAA software. However, recent hardware replacements have only yielded 30-50 percent in application performance improvements as new CPU chip performance begins to level-off in the industry. Investments in hardware alone will not result in NOAA being able to take full use of advancements in HPC. Software engineering is an extremely important part of NOAA's HPC upgrade strategy. This is reflected in national and multiagency initiatives, such as the National Strategic Computing initiative and the Earth Systems Prediction Capability, which not only address high-performance computing architectures, but the software engineering capabilities needed to provide scientific advancement. In order to position the agency to take advantage of the latest, most cost-effective HPC technologies, NOAA is investing $3 million per year in software engineering. The HPC industry is changing and technologies are evolving. In order for NOAA to be positioned to take advantage of these advanced technologies, it must continue its work on re-architecting NOAA's modeling codes. As part of this effort, while simultaneously broadening external engagement, NOAA's Earth Prediction Innovation Center (EPIC) will allow collaboration on model code, which could further improve efficiencies that result in less computational strain. The HPCC program supports EPIC by leveraging its corporate experience to extend existing HPC infrastructure to cloud platforms. HPCC's goal is to provide the best possible cost-effective capacity and architecture for the scientific needs through a mix of public, Federal, and private partnerships. The HPCC program strives to provide access to users of HPC and provide platforms and supporting software infrastructure to widely integrate the scientific community

In FY 2016, NOAA began a 4-year effort to recapitalize two of its R&D HPC systems and to create a sustainable funding profile with which to maintain its existing R&D HPC resources. With the increased funding provided in the FY 2016 and FY 2017 appropriations, NOAA was able to fully replace/recapitalize *Gaea*, its R&D HPC located at Oak Ridge National Laboratory, and will be able to sustain the out-year operations and maintenance of the system. *Theia*, the R&D HPC system located at Fairmont, West Virginia, was recapitalized and replaced by *Hera* with funding provided by the FY 2018 hurricane supplemental. Additional resources provided to NOAA in FY 2017-2019 were used to increase overall R&D HPC computing capacity by adding the *Orion* system through grants awarded to the Mississippi State University. This system is operated by the Mississippi State University.

NOAA's current operational supercomputing acquisition efforts have recently completed awarding a new contract in FY 2020 to enable a two-year transition between the new and legacy contracts. This transition time will allow NOAA to successfully transition operations without interruption to the mission. Supplemental funding provided in FY 2018 and FY2019 will enable NOAA to transition to this new contract while maintaining the same level of computing capacity.

V. BALANCE BETWEEN R&D MODEL DEVELOPMENT AND OPERATIONAL IMPLEMENTATION

Complex earth-atmosphere prediction models provide foundational tools for providing accurate and reliable forecast guidance on various spatial and temporal scales to meet NOAA's core mission requirements. Numerical modeling guidance has been the cornerstone of weather forecasting for decades. Adequate HPC capabilities are central to developing state-of-the-art and next-generation numerical models to conduct research, evaluate, and transition to operations. To meet current and future requirements, NOAA's core enterprise capability and services for computing, post-processing, product generation, real-time dissemination, and data storage needs require continual enhancement and investment. Optimization, modernization, and exploitation of emerging HPC technologies, as well as prioritization of available resources, are critical for the development of the next generation of regional and global models as well as highly reliable operational models.

In order to balance these HPC needs, it is important to understand how NOAA transitions research to operations. Operations and research occur at different time scales that can be depicted as a "funnel" where exploratory research on large development time scales is represented at the top of the funnel and at the bottom of the funnel are the few technology/model advances that actually make it into operations (see Figure 1). For continuous, scientifically sound evolution of operational models and applications, today's research must continuously work on tomorrow's envisioned operational needs. The NOAA modeling enterprise will need to support an integrated approach that facilitates evolutionary as well as revolutionary changes to NOAA's operational production suite of models.

Figure 1: The Research Funnel

Several research to operations (R2O) strategies have recently emerged to facilitate proper planning and prioritization of available resources. NOAA has defined Readiness Levels (RL) to systematically assess the maturity of projects as they transition through the R2O process (Figure 2).[1] Bridging the gap between research and operations is one of the highest priorities in NOAA. To do this, NOAA will need to capture the vast amount of talent and knowledge base residing not only within NOAA, but also within academic, private, and other federal organizations.

[1] NOAA Administrative Order 216-105B: Policy on Research and Development Transitions (http://www.corporateservices.noaa.gov/ames/administrative_orders/chapter_216/216-105B.html)

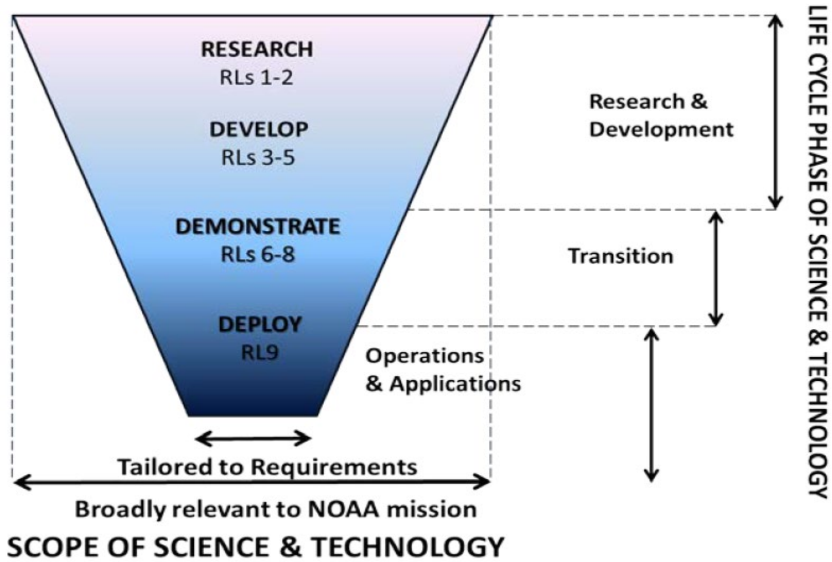

Figure 2: Readiness Levels for Enabling Research to Operations

An example of computational requirements for various NWS Global Forecast System RLs is shown in Figure 3. This shows that RLs 5-6 require almost five times the resources as the operational applications, mainly due to extensive testing and evaluation required to demonstrate the readiness of an improvement for transition to operations. Evidence-based decision-making is critical to maintain the reliability, robustness, and accuracy of operational forecasts. Re-analysis and reforecast experiments are required for model output product calibration and for bias correction. A long history of consistent re-analysis and reforecast data is also required for many downstream applications like hydrology models to simulate rare events (such as drought and flood forecasts) and to support advanced research needs at RLs 1-4. Most of the re-engineering and systematic testing and evaluation happens before handing off the codes for implementation during RLs 8-9. To balance the requirements for research and operations, careful planning is required in allocating R&D HPC for research at various RLs. The allocation of responsibilities for task performance and funding are allocated between the research and operational components and are accompanied by rigorous communication between the parties from proof of concept through adoption into operations.

NOAA's R&D goals directly impact its Weather-Ready Nation strategic goals by improving the science behind forecasts for extreme and severe weather events; developing model improvements, computing, and observations; and including social science to improve risk communication to the public. Additionally, the National Integrated Drought Information System (or NIDIS) Reauthorization Act of 2018 directs NOAA to establish the Earth Prediction Innovation Center (EPIC) to accelerate community-developed scientific and technological enhancements into the operational applications for NOAA's numerical weather prediction (NWP). EPIC will remove barriers in improving numerical weather prediction by enabling scientists and engineers to effectively collaborate in pursuit of improving operational global modeling efforts around a community global modeling system.

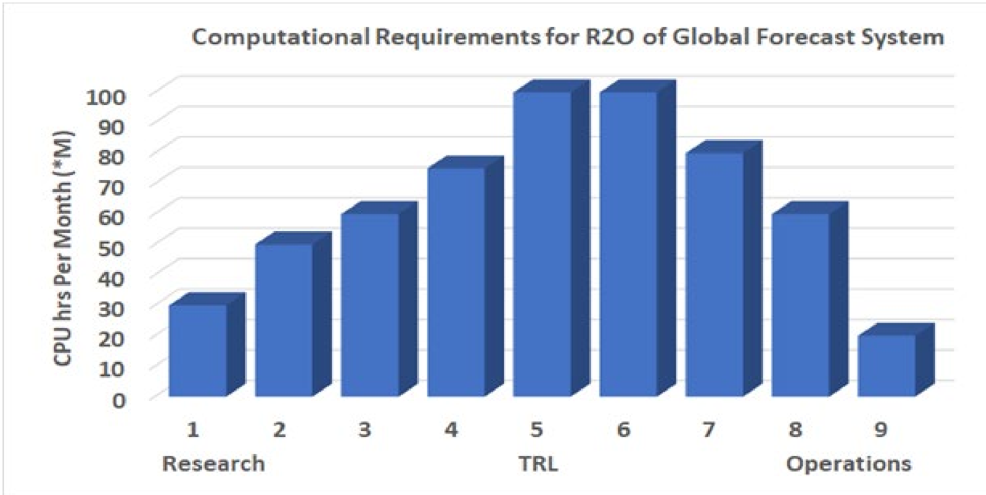

Figure 3: An Example of Computational Requirements for R2O and R&D for Upgrading the NWS Operational Global Model

A. Projecting NOAA HPC Requirements

In order to plan and balance resources, NOAA continually re-evaluates its understanding of HPC requirements to meet mission needs. Recently, NOAA has generated an estimate for an ideal solution for executing the operational production suite over the next 5-10 years, irrespective of available resources, covering core weather applications as well as the growing requirement for generating guidance for related and emerging applications (marine, waves, surge, hydrology, air-quality, space weather, etc.). This analysis can be used both for addressing operational production suite resources, and prioritizing elements in the production suite based on cost-benefit analyses. Projected requirements are based on resolution (horizontal and vertical), length of forecast (typically in hours), number of cycles (model runs) per day, and complexity associated with advanced physics, coupling to Earth system components, and data assimilation techniques. This analysis is instructive to understand relative prioritization in constructing an operational modeling suite that is executable within the resources available to NOAA's HPC Program.

B. Considerations in Planning Computing Needs

NOAA recognizes that expanding and modernizing its forecast applications and products to meet stakeholder requirements requires interaction with, and the support of, the larger research and academic communities. HPC is a vital component for any world-class research program. With the depth and breadth of science areas researched in NOAA, the future demand for HPC resources outweighs the current supply of NOAA's capabilities by 32X. Effective partnerships with the research community require sustained support for research and adequate HPC resources for conducting advanced R2O.

HPC to develop the NOAA production suite should first and foremost concentrate on R2O transition needs. Traditionally, the R2O work has been done on the identical architecture as used for operations to avoid differences related to compute architecture, and to assure stability, minimal code transition, and optimal code execution on the actual operational hardware. More recently, NOAA has been using different and even mixed hardware architectures with success. This includes demonstrating and optimizing the use of cloud-based opportunities.

In general, R2O compute should be significantly greater than the size and capability of compute used in operations. This allows for running several real-time parallel runs, as well as systematic offline development and testing. For data assimilation, the needs may be higher if, for example, the impact of many different instruments is to be assessed simultaneously, as exemplified in Figure 3. Ideally, the greater R2O compute should be relative to the next generation operational machine to assure that additional resources on a new operational machine can be used shortly after the acceptance of the hardware.

C. Prioritization of Operational Computing for R&D

NOAA continually reviews its modeling efforts to make the operational production suite of models as efficient as possible. Most recently, following the recommendations from the University Corporation for Atmospheric Research Community Advisory Committee for National Centers for Environmental Prediction (NCEP) Model Advisory Committee, NOAA has engaged in a major initiative for developing a Unified Forecast System (UFS). This system is an integrated modeling system that unifies scales from individual thunderstorm-scale weather (hourly outlook) to seasonal climate (1-year outlook), and integrates weather, oceans, land, ice, hydrology, and other components in a scientifically sound and economically justifiable way to most efficiently support NOAA's operational mission. The Next Generation Global Prediction System (NGGPS) has adopted a highly scalable Finite Volume Cube Sphere (FV3) dynamic core originally developed at the NOAA Geophysical Fluid Dynamics Laboratory to unify all atmospheric prediction systems within the operational production suite of models. Resources in support of comprehensive water prediction have been essential for the implementation and evolution of the National Water Model and National Ocean Service models, in order to provide the necessary capacity to deliver coupled freshwater-coastal estuary forecasts of total water level in the coastal zone. NWS has taken a lead in developing a 3-year Strategic Implementation Plan that is a long-term Strategic Vision and Roadmap document for the development of the UFS as a community modeling system for research and operations. Ultimately, freshwater and coastal models will be incorporated in the UFS. The development of the UFS will be accelerated through EPIC.

The UFS development requires dedicated HPC resources for a systematic transition of FV3-based numerical modeling systems into operations. At the same time, there is a need for continuous maintenance and upgrades of existing operational modeling systems to account for emerging forecast needs and to accommodate assimilation of new datasets like those coming from GOES-16 and NOAA-20 satellites.

Typically, each forecast system in operations goes through a fixed upgrade schedule. Prioritization of model upgrades depends on various factors including technical readiness of model improvements, mission benefits associated with system upgrades, availability of computational resources in the production window in which the system operates in real-time, upstream/downstream dependencies, and availability of personnel to effect transition to operations. Schedules for these model upgrades are well coordinated among various partners within NOAA and are tracked to address any issues with scope, resources, performance, or risks. These priorities are also determined based on emerging requirements, and incorporation of new observational datasets, where appropriate.

WCOSS development resources are currently reserved for high-priority projects including the development of the NGGPS, systems for predicting severe weather, ensembles to quantify certainty, and annual hurricane model upgrades. Access to the back-up production machine is restricted to only a few developers, hence the use of any available compute resources is limited to those who are closely associated with maintaining the operational models. With new computing capacity added to WCOSS in January 2018, additional computing was made available on WCOSS that significantly increased capacity for 1-3 years. This additional capacity is being used to accelerate development of next generation global and regional prediction systems. These resources will also help disseminate experimental real-time forecast products to the field using NCEP's infrastructure.

NOAA can also achieve strategic enhancement to operational computing through retiring legacy models from operations and replacing them with the next generation global and regional models. This will not only reduce the number of models that need operational maintenance, but also allow human and computer resources to be directed towards accelerating the development of UFS. Use of modern and emerging technologies is critical to enhancing the accuracy and reliability of operational numerical models. Continuous investment in disk and storage space, increased bandwidth and networking, and effective utilization of available resources on the operational HPC are essential for optimal use of computing to support operations.

D. Prioritization of Research Computing for R&D

While many of the numerical prediction models are developed and maintained by NCEP, several other development organizations contribute to the operational suite of models including, but not limited to, the NOAA research laboratories, National Water Center, and National Ocean Service Center for Operational and Oceanographic Products and Services. The broader research community (beyond NOAA) must also be engaged in improving current operational modeling systems and developing next generation unified global and regional modeling systems. Both R2O and operations to research (O2R) are essential in accelerating transition of advanced research into operations.

While HPC resources will always be constrained, there are a few ways to prioritize the R&D resources to support advanced research activities relevant to operations. High priority projects like NGGPS involve both internal NOAA research and the external research community (academic and private), and prioritize research that has reached RL 4-5. Transition plans help HPC portfolio managers prioritize the research requirements and allocation of resources on NOAA's R&D and operational HPC systems. For example, NOAA conducts a demonstration of hurricane model enhancements each hurricane season. The purpose of this demonstration is to evaluate the strengths of promising new approaches that are testable only with enhanced computing capabilities. NOAA evaluates the progress of development work each off-season to identify techniques that appear particularly promising to operational forecasters and/or modelers. Operational programs then blend potential advances into the operational implementation plans through subsequent operational-ready activities or developed further outside of operations within research and development efforts.

Other successful real-time demonstration projects are tested within the NOAA testbeds with researchers and forecasters working side-by-side, including the National Hurricane Center's Joint Hurricane Testbed, the Arctic Testbed, Aviation Weather Testbed, Climate Testbed, Coastal Ocean and Modeling Testbed, Developmental Testbed Center (DTC); Hazardous Weather Testbed; Space Weather Prediction Testbed; and NOAA's various Proving Ground testbeds. In addition to the Testbeds and Proving Grounds, NOAA develops other technologies and techniques relative to: 1) satellites in the Joint Center for Satellite Data Assimilation (JCSDA); 2) climate modeling through the Climate Prediction Center's National Multi-Model Ensemble project; and 3) high-impact and severe weather forecasting through the Earth System Research Laboratory/Global Systems Division's High Resolution Rapid Refresh experimental model program. These focused real-time demonstration capabilities prioritized through various testbeds and proving grounds[2] are expected to advance research into RLs 6-7, which can be handed off to corresponding operational model developers for potential transition activities.

VI. EARLY RELEASE AND COMMUNITY ENGAGEMENT WITH MODELS

NOAA is creating EPIC to accelerate community-developed scientific and technological enhancements into the operational applications for NWP. EPIC will advance the adoption of a unified modeling strategy and establish the UFS in partnership with the scientific community to employ common model architecture and infrastructure co-developed with other Federal agencies. EPIC will require investments in software engineering, software infrastructure, user support services and the ability to run the system on cloud-based, high-performance computing. These investments will be supported through the most efficient and cost-effective vehicle that will best leverage contracts, grants, partnerships, industry and university engagements, and in-kind contributions throughout the scientific community. NOAA, through EPIC, will also provide a management function to communicate, adhere to, and guide the research into operations and operations into research (R2O2R) transition process.

EPIC will provide a framework to accelerate R2O through community-developed enhancements into the UFS and its operational applications. The community includes NOAA, other Federal

---

[2] http://www.testbeds.noaa.gov/

partners (e.g., National Aeronautics and Space Administration, Department of Defense, JCSDA), the broader research and academic community at-large, such as the National Center for Atmospheric Research (NCAR), and the private sector. Only with appropriate contributions from the entire national modeling community will American researchers be able to maintain global leadership in numerical weather prediction. The definition of "community" is important, and not all community efforts will be identical. NOAA is learning from prior and ongoing community modeling efforts (i.e., the Weather Research and Forecasting Model (WRF), Community Earth System Model, Wave Watch Model, Modular Ocean Model (MOM6), etc.) and applying best practices that best meet our specific situation. For example, NOAA, with significant contributions from other Federal agencies and academic partners, has developed a Strategic Implementation Plan for NGGPS. This plan is updated collectively on an annual basis, and applies to broad community engagement in NGGPS development activities for a rolling 3-year period.

The unified modeling system is being built to support the needs of both operations and research. Without that linkage, the research community will be less likely to participate in helping make improvements that will benefit operations, nor will operations efficiently implement research advances. Building a community model involves give and take from both the operational and research sides. Lessons learned, such as from the DTC, have shown us that the community will expect sufficient training, full support (including help desk), and acceptance of scientific advances that will help NOAA build this modeling system. As such, we need to start early to build that infrastructure into the unified modeling system.

Community organization: Different layers of community partners will be established with varying levels of engagement:

- ***Researchers*** should be engaged through Announcements of Opportunity (AO) in order to increase the human capacity needed for long-term improvements (i.e., research funding supports students who will become skilled in the unified modeling technology and environment). For example, AOs from the Hurricane Forecast Improvement Program tripled the research community involvement in development of the Hurricane WRF model.
- ***Core development partners*** (that regularly make substantial contributions) will be granted different roles and access than users that may run the model but not typically directly or actively contribute to its development.
- ***Trusted users*** may be established as a special, limited category that allow greater, early access than normal research users to, in order to conduct early beta testing on the next model version still under development but not yet released to the full community.
- ***Operations***, due to its constraints on reliability, timeliness, and security, will require a unique operational version of the modeling system. A significant goal will be to ensure that the overarching modeling system, while having different variants for research and operations, will have a consistent architecture and infrastructure that will allow improvements made on the research side to be smoothly transitioned into operations.

Also, while not directly contributing to the code development itself, stakeholders can still convey their needs for the predictions to be produced by the modeling system. Therefore, so long as the stakeholder needs fit within the overall mission space of the core development partners, they can help drive the direction of development, resources allocations and prioritization. This could also include a role for basic research partners, such as the National Science Foundation.

Subsets of researchers, core development partners, and trusted users (who have government credentials) have early access to the development version of the NGGPS. This repository[3] is hosted at NCEP's Environmental Modeling Center; a snapshot of this repository was released to the public on March 30, 2018. NOAA is working with the community to gather feedback on model performance. Partners in NWS, OAR, and NCAR are working to design and implement a community modeling framework that will allow the larger research communities to have access to and develop the NGGPS model. While this framework is in development, OAR's model codes relevant to the NGGPS and seasonal to subseasonal prediction efforts are open source and available for download. The ocean component (MOM6) is currently open development and is being used as a reference framework for community model development efforts in the making.

VII. ALLOCATION OF NOAA'S HIGH PERFORMANCE COMPUTING SYSTEMS

NOAA's HPC Board provides strategic guidance and oversight for the execution of the entire HPC enterprise and the allocation of these resources.

WCOSS is primarily dedicated to operational weather modeling. WCOSS is comprised of two systems: a primary and a backup. NOAA dedicates the primary system to operational numerical modeling to support the weather, water, and climate forecasts and warnings. When not in operational use, NOAA uses the backup system for development work. The purpose of this work is to transition modeling improvements into the operational weather model. NOAA allocates nearly 100 percent of the operational systems to NWS for weather modeling and weather R&D transition, which translates directly into improvements in weather forecasts and warnings.

The HPC Board Allocation Committee is responsible for allocating NOAA's R&D HPC resources. All of NOAA's Line Offices except Office of Marine and Aviation Operations receive an allocation on one or more of the R&D HPC systems. This annual collaborative planning across NOAA Line Office, R&D, and associated transition to operations efforts is discussed and prioritized.

Table 1 shows the estimated peak performances of the NOAA compute resources. WCOSS represents the operational machine, and its identical backup (peak teraflops per machine). *Hera*, *Gaea*, and *Orion* primarily support the development of weather and climate applications, and the *Jet* systems are primarily focused on hurricane forecast improvement.

[3] https://vlab.ncep.noaa.gov/web/fv3gfs

Table 1. Distribution of Capacity Across Operations and R&D HPC

Machine	Use	Peak Teraflops
WCOSS Primary	Operational Generation of Model Guidance	4,200
WCOSS Backup	Transition to Operations	4,200
Hera	Research and Development (including Fine-Grain computer)	4,750
Gaea	Research and Development / Long-Term Predictions & Projections	5,290
Jet machines	Research and Development, real-time R2O / Hurricane (HFIP)	1,795
Orion (Mississippi State University)	Research and Development / Weather, Ocean Modeling	5,000

Appendix A: **List of Acronyms**

AOs	Announcements of Opportunity
DTC	Development Testbed Center
EPIC	Earth Prediction Innovation Center
FV3	Finite-Volume Cube-Sphere Dynamical Core Model
GFDL	Geophysical Fluid Dynamics Laboratory
HPC	High Performance Computing
HPCC	High Performance Computing and Communication
JCSDA	Joint Center for Satellite, Data and Assimilation
MOM6	Community Earth System Model, Wave Watch Model, Modular Ocean Model
NCAR	National Centers for Atmospheric Research
NCEP	National Centers for Environmental Prediction
NGGPS	Next Generation Global Prediction System
NOAA	National Oceanic and Atmospheric Administration
NWS	National Weather Service
OAR	Oceanic and Atmospheric Research
OCIO	Office of the Chief Information Officer
R&D	Research and Development
R2O	Research to Operations
RLs	Readiness Levels
UFS	Unified Forecast System
WCOSS	Weather and Climate Operational Supercomputing
WRF	Hurricane Weather Research and Forecasting