United States Department of Commerce
National Oceanic and Atmospheric Administration
National Weather Service
Office of Science and Technology Integration
1325 East-West Highway, Silver Spring, MD 20910-3280, USA

Compute and Storage Resource Assessment for the National Weather Service Modeling Program Office

VERSION 1.1
03/24/2023

https://doi.org/10.25923/3qqg-9b74

Editors: Marc Gasbarro[1,6,7], Hendrik Tolman[1]
Main Contributors: Deepthi Achuthavarier[1,6,7], Farida Adimi[1,6,7], Ligia Bernardet[3], Jacob Carley[2], Stylianos Flampouris[1,6,7], Kevin Garrett[1], Youngsun Jung[1], Daryl Kleist[2], Aaron Poyer[1], Vijay Tallapragada[2], Jeffrey Whitaker[4], Yan Xue[1], Xuejin Zhang[5] , Zhan Zhang[2]

U.S. Department of Commerce (DOC)
National Oceanic and Atmospheric Administration (NOAA)
[1]National Weather Service (NWS), Office of Science and Technology Integration (OSTI)
[2]National Weather Service (NWS), National Centers for Environmental Prediction (NCEP), Environmental Modeling Center (EMC)
[3]Oceanic and Atmospheric Research (OAR), Global Systems Laboratory (GSL)
[4]Oceanic and Atmospheric Research (OAR), Physical Sciences Laboratory (PSL)
[5]Oceanic and Atmospheric Research (OAR), Atlantic, Oceanographic and Meteorological Laboratory (AOML)

[6]Science and Technology Corporation
[7]Affiliation when contributions made

Change Record Page

Version	Date	Page affected	Description
1.0	12/01/2022	all	Formal first draft generated and reviewed by OSTI Modeling PO and UFS-R2O Project Leads
1.1	03/24/2023	most	Updated based on feedback from NOAA NWS and OAR Directors; reformatted

Table of Contents

1 Introduction

1.1 Purpose

The Modeling Program Office (STI-M) within the National Weather Service's Office of Science and Technology Integration (NWS/OSTI) assesses the compute and storage needs for various weather applications moving to the operational Production Suite at the National Centers for Environmental Prediction (NCEP). The programs and projects covered in this assessment are those that are managed by STI-M and are in various stages of development, but all have an eye towards operational implementation within the next 2-5 years. Resource constraints are always a concern with any project or program, but no resource constraint within NOAA's numerical weather prediction (NWP) applications presents a higher risk than the need for and lack of compute power and storage capacity. Computational High-Performance Computing (HPC) resources for Operations and Research and Development (R&D) are critical for advancing science and improving operational outcomes for the NWS. This assessment presents estimated compute and storage needs and recommendations necessary to inform NOAA's decision makers on funding gaps and resource needs required to successfully plan and transition STI-M's managed programs and projects from research to operations (R2O) and to help build a Weather Ready Nation.

1.2 Background

The NWS Modeling Program Office manages a host of programs and projects, chiefly the UFS-R2O Project. While each program and project within STI-M have differing needs, all have the goal to accelerate and advance research initiatives to support and accelerate NWS operational model development and improve forecast accuracy. STI-M focuses on more mature technological advancements higher in the readiness level spectrum with the focus to fold the integration of NWP applications into the operational model suite within NCEP Central Operations. These efforts are compute and storage intensive. In line with the 2021 Report on the Priorities for Weather Research (PWR), "Improvements in weather forecasts are directly limited by the availability of sufficient computing resources to develop, test and operate next-generation forecasting technologies." This document estimates the "sufficient computing resources" needed to "develop, test and operate" these technologies.

NOAA's Office of the Chief Information Officer (OCIO) has the ultimate responsibility to ensure a comprehensive strategy for obtaining and implementing enterprise-wide supercomputing support to NOAA's research and operational missions. Program Offices are best equipped to account for the compute and storage needs for their programs and projects and communicate the resources required to OCIO via their respective channels and committees. Sources of on-premises computing resources primarily focus on the Research and Development High Performance Computing (RDHPC) Systems for R&D and on the Weather and Climate Operational Supercomputing System 2 (WCOSS2) for transition to operations testing and operational implementation. Cloud architecture platforms present a growing and more feasible alternative to on-premises computing through a myriad of cloud-service providers (CSPs). Finally, storage is another resource aspect that is primarily available through the High Performance Storage System (HPSS) or through the NOAA Open Data Dissemination (NODD) program, formally known as the Big Data Program.

1.3 Problem Statement

NOAA is vastly under-resourced in compute capacity to fulfill its NWP mission. As weather and climate modeling capabilities increased, so has the insatiable appetite for computing capabilities and corresponding infrastructure to support data transport and storage. According to the PWR, a “four-fold increase in model resolution” would translate to needing “on the order of 100 times the current operational computing capacity” by 2031. Additionally, NOAA severely lacks in R&D and Transition to Operations (T2O) computing capacity. As of early 2022, NOAA’s R&D to Operations ratio is 1:2 whereas other international weather centers have the ratio as 3:1 to 4:1. Despite ongoing efforts by NOAA to procure more resources to bring the ratio above 1:1 in FY23, that is still not enough to fulfill NOAA’s R&D mission areas. This translates downward to STI-M’s programs and projects, particularly UFS-R2O, as there is increasing competition across NOAA’s portfolios for compute and storage resources. Beyond this lack of R&D and T2O computing to support improvement of operations, NOAA in general also lacks resources including HPC to do more fundamental research to feed the R&D and T2O pipelines. The latter, however, is a more structural resource need, and is outside the scope of this assessment.

For UFS-R2O specifically, there currently is not enough resources to operationally implement UFS-R2O’s applications by FY24 and beyond as they, collectively, will exceed the capacity of WCOSS2. When factoring in R&D and T2O needs, computing resource shortfall problems become more acute. For example, new UFS applications are about 3-4 times more expensive than ones used on the current operational systems. Furthermore, for R&D it takes only one instance for some of the forecast systems running at full resolution to take up nearly the entire capacity of UFS-R2O’s current R&D HPC allocations. Adding to the problem is that the compute allocations are spread across numerous discrete RDHPC systems and that UFS-R2O does not have its own dedicated HPC resources; therefore, relying on “in-kind” contributions from other RDHPC portfolios.

STI-M surveyed project leads and program managers to project compute and storage needs. The needs presented are estimates based on the projected needs of each application according to the expected incorporation date into NCEP's Production Suite. It is important to note that these needs assume full funding availability, no major schedule impacts, overcoming any known significant software and technical risks and infrastructure shortfalls prior to production, avoiding major unknown risks and, of course, adequate availability of computing and storage resources.

2 Compute and Storage Resources

2.1 Programs and Weather Applications

STI-M manages several programs and projects, with the UFS-R2O Project as the most comprehensive and resource-intensive. Programs and projects are part of the Modeling Program Office's annual budget process, and include:

Air Quality: Supports the National Air Quality Forecast Capability (NAQFC) to provide operational air quality forecast guidance and improve the basis of air quality alerts and provide air quality information. The air quality forecast is the collaboration between research and operations at NOAA and external partners. The program supports the development and improvement of Air Quality models, the online-CMAQ at 13-km and 3-km resolutions over the North American domain and the GEFS-GOCART coupled meteorology-aerosol forecast system.

COASTAL Act: The Consumer Option for an Alternative System to Allocate Losses (COASTAL) Act requires NOAA to produce detailed “post-storm assessments'' in the aftermath of a damaging tropical cyclone that strikes the U.S. or its territories. Using output from a hindcast model (termed the “Named Storm Event Model'' (NSEM) by the Act), the assessments indicate the strength and timing of damaging winds and water at a given location in the area impacted by the tropical cyclone. The Act further requires NOAA to create a “Coastal Wind and Water Event Database'' (CWWED) for public access.

HFIP: Hurricane Forecast Improvement Program (HFIP) aims to improve guidance for hurricane track, intensity, and storm surge forecasts, with an emphasis on rapid intensity changes. Established under the Weather Act 2017, HFIP continues to advance through the development of the Hurricane Analysis and Forecasting System (HAFS). HFIP also includes the HFIP Real-time Experiment (HREx). Primarily during Hurricane season, the HREx conducts potential operational configurations of the experimental HAFS that are evaluated at the HFIP Real-time Experiment for future operational model upgrade considerations.

Weeks 3/4: The Weeks 3-4 Program aims to improve extended-range weather outlooks by extending numerical model guidance out in time, building reforecast and reanalysis capabilities, innovating forecast products, and supporting applied research in the sub-seasonal forecast range. The program includes modeling improvements for the GEFSv13, Weeks 3-4 products improvement, Arctic Sea Ice Prediction and Flash Drought Prediction, and supporting the Climate Prediction Center (CPC) with post-processing and improvement of forecast products.

UFS-R2O Project: The Unified Forecast System - Research to Operations (UFS-R2O) is a broad collaborative project between the National Weather Service (NWS) and non-NWS researchers created for efficiently incorporating cutting-edge research and innovation into the NWS operational forecasting systems. Its primary goal of Phase I is to deliver three new operational systems based on the UFS: A global, coupled ensemble prediction system for medium-range and sub-seasonal prediction, a convective scale regional ensemble prediction system for severe weather over the North American domain, and a storm-following Hurricane Analysis and Forecast System (HAFS) for high resolution tropical cyclone predictions across the globe. Specifically, Phase 1 of UFS-R2O is broken down into three Application Teams (ATs) with four modeling suites and six cross-cutting teams (CCTs). In addition, there are two additional modeling suites currently considered for Phase 2 (beginning 4QFY23) included here for planning purposes. As each modeling suite is detailed below, it is important to note that this is a very-high level description of each application. It is also important to note that every application requires Data Assimilation (DA) and post-processing. Compute and Storage requirement estimates will follow in section 2.4.

1. **HAFS.** The Hurricane Analysis and Forecast System (HAFS) a multi-scale, multiple moving nested modeling system with a data-assimilation package and ocean-wave coupling. HAFS, established as part of HFIP, falls under the purview of the Hurricane Application Team within UFS-R2O. They develop the next generation hurricane forecast system with the aim to replace the Hurricane Weather Research and Forecasting (HWRF) & Hurricanes in a Multi-scale Ocean-coupled Non-hydrostatic (HMON) models and to develop tropical cyclone initialization, atmosphere and ocean coupling, physics upgrades, and adding telescopic moving nests within HAFS. The HAFSv1 release encompasses two configurations, to replace HMON and HWRF, each with one storm-following 2 km moving nest in a regional parent 6 km domain, 81 vertical layers, a cadence every 6 hours, and a forecast length of 126 hours.

2. **GFSv17/GEFSv13.** The Global Forecast System (GFS) and the Global Ensemble Forecast System (GEFS) are fully-coupled forecast model applications that are part of the Medium Range Weather/Seasonal to Subseasonal (MRW/S2S) AT. UFS-R2O advances forecasting

capabilities on both systems by developing and releasing the latest versions (GFS v17 & GEFS v13) into the operational production suite at NCEP. GFSv17 plans to run on a 13 km horizontal resolution global domain with 127 vertical layers with a cadence every 6-hours and forecast length of 384 hours (16 days); however, pending available coupled model resources, GFSv17 could ideally increase horizontal resolution to 9 km. Though not finalized, GEFSv13 will likely run a 31-ensemble member system on a 25 horizontal resolution global domain with 64 vertical layers with cadence every 6 hours and forecast length of 16 days and cadence every 24 hours out to 35 days. GFS/GEFS will be initialized using the Global Data Assimilation System (GDAS) that includes a weakly coupled ocean-atmosphere-sea ice-land ensemble data assimilation system that is under development.

3. **SFS.** The Seasonal Forecast System (SFS) is another application that is part of the MRW/S2S AT. It will likely be developed in Phase 2 of the project. The UFS-R2O will develop and release a first version of a fully-coupled SFS which will provide probabilistic subseasonal to seasonal (S2S) forecast guidance going from a period of 4 weeks to an end goal of 2 years. The SFS configuration is still under consideration; however, the goal is to release the first version of SFS with a 25 km horizontal resolution. Exact cadence and resolutions (vertical and horizontal) will depend on the forecast length period. For example, for the subseasonal prediction with a forecast length of 45 days, a cadence of 24 hours may be appropriate. For longer-time frames (seasonal to seasonal) with a forecast length of 360 days, a cadence may only be required once a week to once per month. Ensemble sizes are also under consideration and will vary depending upon the length of forecast being run. Ensemble size possibilities range from 50 to 200 members. In addition, exact DA for SFS needs are unknown; however, the basic idea is to initialize from the coupled GDAS by leveraging developments from the GFS/GEFS initialization. This document will cover the compute costs of the various possibilities in section 2.4.1.

4. **RRFS.** The Rapid Refresh Forecast System (RRFS) is a convection-allowing ensemble forecast system that is underpinned by the Short-Range Weather/Convection Allowing Model (SRW/CAM) application. UFS-R2O advances short-range forecasting capabilities by developing the SRW/CAM application. The RRFS will run a limited area model on a North-American domain with lateral boundary conditions provided by the GFS. It will simplify the operational product suite by replacing several legacy regional models including the North American Mesoscale (NAM), Rapid Refresh (RAP), Short-Range Ensemble Forecast (SREF), regional NAM nests, High-Resolution Rapid Refresh (HRRR), and High-Resolution Ensemble Forecast (HREF) forecast systems. It will run with a 3 km horizontal grid-spacing, 65 vertical layers, and feature hourly updates with a hybrid ensemble-variational data assimilation system. Forecasts are issued for the full North American domain every hour out to 18 hours and every 6 hours to 60 hours. Forecasts to 60 hours also include a 10

member ensemble forecast that is initialized off a subset of the RRFS's data assimilation system ensemble.

5. **3D-RTMA.** The 3 Dimensional - Real Time Mesoscale Analysis (3D-RTMA) is a high-resolution assimilation system developed by the SRW/CAM Application Team (AT) that will present real-time, rapid-updated analysis of 3D atmospheric fields. It includes both the real-time component (3D-RTMA) and its delayed real-time companion system, UnRestricted Mesoscale Analysis (URMA). The baseline version of 3D-RTMA will run on a 2.5 km horizontal grid-spacing, with a cadence of every hour that will later extend to every 15 minutes. It will utilize the RRFS forecast as the background as well as the RRFS ensemble data assimilation members for flow dependent error covariances.

6. **WoFS.** The final system that is envisaged to be supported by the SRW/CAM Application is the Warn on Forecast System (WoFS). It is considered for development during Phase 2 of the project. The WoFS will be a convection-allowing model with a weather-adaptive regional and relocatable domain (~ 900 x 900 km) run with an output cadence of at least every 15 minutes with a goal of having forecast length of 6 hours. It will include an ensemble forecast system of 18 members providing 6 hour forecasts every 30 minutes with initial and lateral boundary conditions from the RRFS ensemble. It aims to include 36-members in its ensemble data assimilation system, likely leveraging components of the RRFS's DA ensemble system. Similar to the SFS, final configurations are still under consideration; however, horizontal resolutions will likely extend down to 1.5 to 2.5 km.

7. **CCTs.** UFS-R2O has six cross-cutting teams (CCTs) that all contribute to the development of the applications mentioned above. The six CCTs are Atmospheric Composition, Atmospheric Physics and Dynamics, Data Assimilation and Reanalysis & Reforecast, Marine Components, Modeling Infrastructure, and Verification and Post-Processing. Each of these teams cross-cut their work interchangeably with the ATs. As such, the vast majority of compute/storage resources the CCTs require are included in the estimates for the applications in section 2.4.; therefore, this document will not duplicate CCT needs.

8. **Other ATs.** Additional ATs are possible in a follow-on phase of the UFS-R2O Project to include applications in air quality, coastal weather, hydrology, and space weather. Future iterations of this document may include compute and storage costs for additional applications.

2.2 Weather Modeling Costs and Factors

High Performance Computing (HPC) and storage infrastructure needs will vary considerably depending on numerous cost factors and decisions on capabilities and tradeoffs between increased needs and finite resources.

There are three main elements and numerous cost factors that affect compute needs. In modeling, the three main elements of compute costs are data assimilation, forecast model (and ensemble) runs, and post-processing/product generation. The first element, DA, has computations that include the processing of observations and forward calculations, computation of increments, and the minimization iteration process. Costs for all elements except for observation processing depend on the resolution of the model background and the analysis and the number of ensemble members for ensemble DA. The second factor, forecasts, is typically the costliest and is covered in the following paragraph. Compute numbers listed in this plan consider the best estimates for DA and forecasts. The last element, post processing, is typically minimal from a compute perspective, but may require significant storage, and access to model output and observations. Post processing may include product downscaling and visualization considerations. Post-processing may become more intensive if model applications become input-output (IO) limited in terms of read/write from disk to memory. Setting up a more heterogeneous system may help in terms of post-processing costs.

When considering computational cost factors, especially forecasts, resolution (horizontal and vertical) is the most significant, followed (in no particular order) by cadence, domain size, forecast time steps, length of model run, and number of ensemble members. These factors are the largest and most quantifiable and, as such, project engineers and managers can approximate the compute needs from them. For example, increasing the horizontal resolution by a factor 2 along with a likely change of time step of a factor of 2 increases computational costs by a factor of 8. Additionally, for example, doubling the vertical resolution alone increases the computational cost by 2, assuming there are no additional time steps needed for vertical resolution increases. Less quantifiable, yet still important, are factors such as code optimization, coupling efficiencies, IO interference between applications, graphics processing unit (GPU) processing, parallel computing scalability, and physical parameterization. IO interference, in particular, is increasingly becoming a bottleneck in transferring data and model output. Finally, there are coupling modeling components. Although models have multiple components coupled together (ice, ocean, waves, etc.), it is the atmosphere component and dycore that is, by far, the largest computational cost.

Prior to implementation, project engineers, managers, and developers need to consider a scaling factor for R&D and T2O costs. This is to test each system extensively and to tune the system for optimal forecast performance. The scaling factor is the ratio between R&D/T2O resources and operational resources needed. It accounts for running the baseline model plus real time retrospective simulations, multiple parallel runs (testing and experimentation), regression testing, configuration changes, tuning/optimization, and unknowns that arise during the development process. T2O also accounts for pre-implementation testing by running the model applications on the same hardware that will be used for operations (mainly WCOSS2) in order to test code reliability, ensure operational environment compatibility, and verify proper configuration management run on the same hardware. While most weather modeling projects use a scaling factor anywhere from 3 to 6 for compute costs, a good assumption from the environmental modeling community targets a factor of 5 and this is the factor that UFS-R2O Project Leads, among others, are utilizing to estimate R&D and T2O resource needs across applications.

In addition to above, Reanalysis and Reforecast (R&R) requires a special consideration. R&R is an integral part of NWP operational implementation, as both are necessary for end-users to calibrate and bias-correct numerical guidance. Reanalysis provides a consistent, gridded record of weather and climate by assimilating historical observations into a modern weather forecast model for use as initial conditions for a historical reforecast. Reforecasts are critical for correcting model biases, estimating ensemble skill, calibrating real-time operational forecasts, as well as training machine learning algorithms. For consistency, developers strive to use the same model configurations for both reforecasts and reanalysis.

Another consideration is Code Retirement. As new models deploy into the production suite, legacy products are turned off by NOAA Central Operations (NCO); however, there may be a small burn-in period to ensure a smooth rollout into operations and streamline upstream and downstream dependencies all while doing no harm to operators and the public. The most impactful implementation is RRFS as it will replace a myriad of products. Note that whereas, for instance, the RRFS implementation will result in retirement of many older models, the production suite is also including many legacy products, mostly old product formats and layouts that are mimicked by newer models to assure that programs (e.g. AWIPS) keep working after legacy models are retired. Many such products have accumulated over time, and whereas they generally require only minor compute resources, their operations and maintenance (O&M) and transition to new machines is highly labor intensive. Hence, this is a second needed focal point of Code Retirement, that is mostly out of scope for the present assessment.

Finally, there is the storage consideration. There are two main types of storage, disk and archived. Disk storage entails temporarily storing instances for the model runs and observations. These instances are eventually overwritten by subsequent model runs; however, high disk usage can affect memory usage and HPC performance. Furthermore, enough disk storage is necessary to incorporate previous runs in subsequent model runs and for forecast verification purposes. Archived (or deep) storage saves the model runs permanently. Models get archived and consume large quantities of space and present the larger storage consideration for costs. For storage, factors such as the number of output parameters, changes in output frequency (cadence), and efficiency of data compression affect storage usage.

2.3 Timeline

The timeline for implementation of UFS-R2O applications into the Production Suite is not final and is subject to changes. This document assumes the latest estimated milestones for implementation. It is important to note that production delays may continue to occur given budget and compute resource uncertainties and other programmatic and/or technical high-impact risks to applications. Table 1 lists the scheduled timeline (subject to change) for the implementation completion of each application's next (or first) version release. Of note, all of the applications listed, other than SFS and WoFS, have follow-on implementations notionally planned approximately 2 years after the next release.

*Table 1: Timeline for completion of UFS-R2O Application Releases (as of February 2023) **

**Subject to changes*

HAFSv1	**RRFSv1**	**3D-RTMA / URMAv3**	**GFSv17 / GEFSv13**	**SFSv1**	**WoFSv1**
Q4FY23	Q1FY25	Q1FY25	Q3FY25	2027-2028	Q1FY27

2.4 Gap Analysis

2.4.1 Compute and Storage needs

Compute and storage needs listed here are the best estimates for the programs and applications. The needs listed are taken from multiple sources to include program managers and project leads, program and project documentation, and calculations from OSTI HPC/compute expertise (also briefed at the 102nd AMS Annual Meeting). They are, at best, ballpark estimates, especially for the R&D and T2O needs. This is not to be authoritative nor is it final as needs change when scope and direction of programs and projects change. Table 2 highlights best-estimates for STI-M programs in terms of central processing unit (CPU) hours and storage (approximate average needs per year from FY 23-26) for R&D and T2O. It is important to note that the Weeks 3/4, HFIP, and AQ programs include some efforts that are undertaken within UFS-R2O; therefore, the estimates in Table 2 only include resources needed for the non-UFS-R2O portion.

Table 2: STI-M Program Compute and Storage estimates for FY 23-26 (average p/yr)
** Only includes estimates for efforts outside of UFS-R2O*

	Cores hrs (Million)/month	Disk Storage (TB)	Archived Storage (TB)
Air Quality (AQ)*	4-5 M	200-400	5-8
HFIP*	45-50 M	1100-1500	1500-2500
Weeks 3/4*	1.6M	100	500
COASTAL Act	0	0	400
Total	**50.6 - 56.6M**	**1400 - 2000 TB**	**2400 - 3400 TB**

Table 3 contains UFS-R2O best-estimates to include implementation, R&D, T2O, and storage. It is important to mention some assumptions. First, for R&D and T2O, a scaling factor of 5 is used, as explained in Section 2.2. Additionally, the R&D and T2O needs will continue for all applications after the next release as the ATs are developing and testing for follow-on releases; however, Table 3 primarily focuses on those needs for the next release. Second, the Petaflop (PF) estimates given in Table 3 represent the needed compute peak performance assuming the work is performed with a constant resource allocation 24x7, even if the cadence is much larger than one day (for example, the SFS). This is done purely for consistency and comparison. It is worthy to note, that for large cadences and/or using cloud resources, CPU hours or node hours may be more relevant. Third, the compute estimates represent the high-water mark for each application. This presents the maximum load that each application would utilize. Fourth, the resource needs in Cores and PF are calculated based on WCOSS2 specifications and utilization. Even though a significant amount of R&D will be on the RDHPC systems and/or cloud, it's important to keep a consistent comparison to T2O testing and operational implementation loads on WCOSS2. Finally, the disk storage listed is based on the new amount needed for that particular cycle and does not include storage of previous runs that is kept on disk for initialing subsequent model runs and for forecast verification purposes. The amount needed depends on the uses of each application, but will likely be substantially larger than what is required for just one model run. It is also important to note that Table 3 does not consider the amount of disk storage needed for the R&D and T2O aspects, which alone needs well over 10 PB per year.

Table 3: UFS-R2O Project Compute and Storage estimates
** Operational Configurations not solidified; estimates are rough*
***For SFS, cores unknown; PFs better measurement of needs due to infrequent cadence & unknown configuration*
****For SFS, compute & storage amounts "assume" cadence of 24-hours; reducing cadence reduces estimates shown*

	Operational Implementation Estimates				R&D/T2O
UFS/R2O Application	Cores (K)	Compute (PF)	Storage (TB)		Compute (PF)
			Disk (p/ cycle)	Archived (p/ day)	
HAFS	77	2.6	0.4	1.6	13
GFSv17/GEFSv13	110	4.0	134	36	19.8
*SFS	**unknown	***5 - 90	200 - 2000	***200 - 2000	***25 - 450
RRFS	184	6.6	20	110	33.1
3D-RTMA	21	0.75	0.04	3.2	3.8
*WoFS	5 - 30	0.15 - 1.05	< 0.001	0.013 - 0.017	0.8 - 5.25
Total (minus SFS)	**~ 392 - 417 K**	**~ 14.1 - 15 PF**	**~ 155 TB**	**~ 151 TB**	**~ 70.5 - 75 PF**

For SFS and WoFS, Table 3 highlights a range of numbers based on potential model configuration. Since implementation is not scheduled for both SFS and WoFS until 2027 or later, exact configurations are not yet solidified. For WoFS, unknown configurations are primarily focused on the horizontal resolution and ensemble use and size. WoFS estimates in Table 3 include running its own 18-member forecast ensemble, utilizing RRFS's DA ensemble, assuming 70 vertical layers, and a range of possibilities from 3 km resolution (low-end numbers) to 1.5 km (high-end estimates). Table 3 WoFS estimates assume domain size and cadence as mentioned in section 2.1.

The SFS is more complicated. There are several possibilities for SFS as developed from an OSTI tool and briefed at the 102nd AMS Annual Meeting. This assessment assumes that SFS may run two separate models, one for subseasonal forecasts out to 45-days and another for inter-seasonal forecasts out to 360-days. Both subseasonal (SSFS) to seasonal (SFS) forecast solutions assume 128 vertical layers. The main parameters left are cadence, horizontal resolution, and ensemble sizes. Controlling for cadence, Tables 4 and 5 illustrates possible tradeoffs in compute size for both models. Compute estimates are extremely rough and, therefore, a range is summed in Table 3. When looking at the estimates, consider that resolution becomes less important than ensemble member size for the longer-range SFS as ensemble sizes drive accuracy at very long ranges vs. resolution driving accuracy at much shorter time spans. Only for consistency and comparison purposes with other applications, the estimates given in Tables 4 and 5 are for a cadence of 24-hours. A 24-hour cadence is unrealistic and likely unnecessary, especially for seasonal time frames and, therefore, exact compute costs will likely be much less than presented. Decision makers, for example, may choose to run the SSFS variant 1-2x per week and the SFS variant 1-2x per month. This will greatly reduce compute costs. In addition, the infrequent nature of its operation will necessitate gravitating away from WCOSS2's use and, likely, towards a cloud architecture option, even for operations. Finally, project leads should measure the final configuration needs in core/CPU hours per month or cores per model run as that statistic will highlight a more realistic figure than just PF.

Table 4: SubSeasonal Forecast System (SSFS) Estimates (in PF) - 45-day length

Ensemble member size (top row) vs. horizontal resolution (left column)

	31 members	50 members	100 members
35 km	0.71 PF	1.14 PF	2.27 PF
25 km	1.93 PF	3.12 PF	6.23 PF

Table 5: Seasonal Forecast System (SFS) Estimates (in PF) - 360-day length

Ensemble member size (top row) vs. horizontal resolution (left column)

	28 members	50 members	100 members	200 members
50 km	0.2 PF	2.9 PF	5.8 PF	11.6 PF
35 km	4.8 PF	8.5 PF	17.0 PF	33.9 PF
25 km	13.0 PF	23.3 PF	46.5 PF	93.0 PF

A critical additional need is R&R. The advancement and production of R&R requires large amounts of dedicated resources for observation data management and considerable computational resources. Under-resourced reforecasts can delay operational model implementation. NOAA, unlike the ECMWF, normally runs their R&R on non-operational platforms; therefore, no WCOSS2 resources are taken.

A 30-year reanalysis system for GFS/GEFS as well as for the SFS (0.25 degree resolution, 80-ensemble member, coupled DA) cost estimate is approximately 1.42B CPU hours. This would equate to running 216K cores 24/7 for 9 months. A trade-off may be to halve the resolution to 0.50 degree as this would reduce the reanalysis cost to 677M CPU hours. The reforecast costs are an additional 355M core hours. This brings R&R estimated total to 1.775B core hours, if utilizing the higher 0.25 degree reanalysis resolution. In addition, R&R may be required for SFS pre-implementation. If required, it would present an additional cost. A lesser-cost option for R&R is presented in the Recommendations section.

Finally, storage for legacy code is considered. As mentioned above, ideally, legacy models are retired as new ones are implemented. The biggest upheaval will be RRFS since it replaces a suite of models mentioned above. If legacy code must be maintained after RRFS implementation, that will cost 1963 GB of archived storage per day.

2.4.2 Limitations and Impacts

All applications are limited in terms of the amount of compute available on WCOSS2 for model implementation and RDHPC Systems for R&D. Table 6 lists the capacity limits for WCOSS2 and the combination of the four main RDHPC systems along with the HPSS for archived storage. It is important to note that RDHPC and HPSS systems are currently running at near full capacity from users across NOAA. Furthermore, the aggregate amount of compute power listed in Table 6 is for reference purposes only, as the problem is acutely compounded by developers' time consumption in migrating and porting data and code between the four discrete RDHPC systems. The fragmented nature of the RDHPC systems makes effective use of resources less manageable than if they were a monolithic system. WCOSS2 includes 12.2 PF, of which 11.8 PF is for compute (327,800 cores) and 0.3 PF for big-memory (8,448 cores), which is needed for pre/post-processing actions. WCOSS2 also has a back-up system that, when available, NCEP uses for R&D and T2O live parallel test runs. The resources needed (as listed in Tables 2 through 5) exceed the current capacity of on-premises resources (as listed in Table 6).

Table 6: Advertised Capacities for NOAA's Primary HPCs

** WCOSS2 back-up system has equal amount of resources for potential T2O/R&D use*

***CAVEAT: This number is aggregated among RDHPC Systems for reference only*

	***WCOSS2**	**RDHPC Systems** **(Gaea/Hera/Jet/Orion)**	**HPSS**
Performance (PF)	12.1	**17.7	-
Compute (M Core hrs/mo)	236M	246M	-
Storage (Disk or Archived) (PB)	26.2	49.3	310

Significant impacts will occur without the required amount of compute and storage resources. The biggest impact will be to the schedule, for without the resources, deployment timelines to the NCO will slip from the dates presented in Table 1. This includes not having the resources to perform the testing and tuning required to get the performance of the new systems to the point that they exceed the current operational systems and to generate the required supporting reforecast datasets. Other major impacts involve a reduction in technical capabilities, capability tradeoffs, and task prioritizations to reduce the computational cost in order to align with the resources available. These impacts extend to unique needs such as R&R. An ad-hoc R&R needed to meet pre-operational requirements risks inconsistencies with real-time forecast models and DA and, thus, will pose challenges for statistical post-processing techniques. In addition, delaying R&R due to resource constraints risks deployment timelines for GFS/GEFS. There are opportunities for some applications to utilize cloud computing, as is discussed in the following section. Additionally, recommendations to discuss mitigating some impacts are discussed in Section 3.

2.5 Cloud Opportunities

Cloud computing is becoming more efficient, more economical, and more practical. There are advantages and tradeoffs of using cloud computing architecture platforms. A major advantage is the increased continuity, development, and expertise of advancing existing cloud projects such as 3D-RTMA and RRFS on-the-cloud prototypes. Both project prototypes primarily utilize commercial cloud services for compute, resource optimization management (cluster, file system, parallel works, workflow), and storage. Cloud compute yields faster processing compared to on-premises HPC (2021 RRFS on-cloud prototype runs were ~15% faster compared to NOAA RDHPC). Processing time reductions means using fewer instances, thus saving cost. With cloud providers, instances are more available, a huge advantage for testing and development of parallel applications vs. on-premises batch queueing prioritization.

Particularly, when cloud resources are not used, they are banked for later use; whereas, when on-site compute cycles are not used, they are lost forever. This implies a higher use efficiency of cloud resources compared to on-premises resources that need to be considered in cost comparisons. Furthermore, there is an increasing need for scalable operational modeling on-demand compute, which is not feasible with finite on-premises HPC resources. Other cloud service advantages include access to the latest generation of processors, choice of CPUs/GPUs, flexible file systems, infrastructure flexibility, configuration/customization flexibility (e.g. firewalls, instances, load balancers, operating systems, user permissions, etc.), easier data movement and model dissemination, easier access for developers, more efficient throughput, increased uptime, minimal maintenance costs, network interconnect options, scalability, leverage of cutting-edge machine learning (ML) and artificial intelligence (AI) capabilities, and improved resiliency (automatic back-up).

Cloud efforts in NOAA have historically focused on research, development, and testing. This is an ideal avenue as cloud offers more effective and flexible options for developing and running tests and prototypes. It also makes it easier for academic and non-NOAA researchers to access codes and repositories to enable their progress towards collaborative efforts in helping NOAA develop the next-generation of NWP capabilities. OSTI assists NWS sponsors by managing some R&D investments for cloud services through NOAA’s Virtual Lab (VLab). Some programs and projects, such as the Earth Prediction Innovation Center (EPIC) and many developments ongoing within UFS, are building their capabilities around cloud implementation. This consideration is discussed further in the Recommendations section.

There are tradeoffs and considerations relating to costs and impacts to operations. A major tradeoff weighs the cost of on-premises HPC ownership (maintenance and idle resources) vs. the cost for on-demand cloud computing. Cloud costs can be prohibitive for high-egress network traffic or very computationally extensive projects. Cloud costs primarily include compute (e.g. instance size, number of CPUs and GPUs), disk space (e.g. volume size, disk speed), shared storage (e.g. network file systems, object storage), and egress networking. However, CSPs do offer discounting for predictable workloads. In addition, R&D can take advantage of lower spot or reserved pricing vs. higher-cost on-demand ‘spike’ computing, although on-demand pricing continues to trend downward. Other considerations include cloud expertise availability (which NOAA lacks), training, costs to design cloud-ready applications, cloud-migration costs, data egress costs, HPC network interconnectivity (on-premises and cloud), bandwidth, latency, network security, and non-standard observation data restrictions.

Storage infrastructure presents another opportunity to move toward the cloud. Archiving every model run adds up. Program managers and project leads should consider exploiting the NODD for storing model outputs and data. This would eliminate storage and egress fees. It is easy to access and use for both the developers and the community. Also, storage is easier to cost out than compute.

Regardless of the type of cloud computing and storage architecture considered, per the NOAA Cloud Strategy, the use of cloud services is NOAA's "desired future state". This is in-line with recommendations in the PWR to leverage "cloud platforms", explore "cloud architectures", and invest in "cloud computing". NOAA's Cloud Strategy further states, "IT investment decisions should fully and credibly consider cloud as a potential alternative" to traditional on-premises HPC infrastructure; however, program managers, principal investigators, and project engineers should make these choices "based on a deliberate decision process" and craft a cost-benefit analysis for each application that are good candidates for cloud compute.

3 Recommendations

This section provides several general recommendations on increasing computing and storage capabilities. These recommendations include obtaining more R&D HPC resources, exploring cloud computing, improving efficiencies for better compute optimization, resolving IO limitations, exploring alternate technologies, crafting mitigation strategies, analyzing tradeoffs, obtaining resources for R&R, and finding adequate storage solutions. These recommendations, among others, are in-line with the NWS OSTI's 5-year strategic vision to *Advance and Enable the Weather-Ready Nation Vision through the Innovative Application of Science and Technology* by successfully managing and accelerating the transition of research and development to NWS operations.

According to PWR Recommendation FE-6, "The federal government should immediately invest in substantially more computing resources dedicated to weather forecasting research, development, testing and operations". Additionally, the report states that "major investments are required in computing resources, including cloud computing, next generation computers, storage, and bandwidth, especially for research computing, but also for the operational implementation of more comprehensive models". In determining the best solution(s) for computing and storage, program managers and application teams should weigh different cloud options. Program managers and project leads should plan future application designs with an eye toward utilizing cloud architectures and networking, along with assessing the true cost burden for on-cloud use by identifying the project's most cost dominant aspects in order to achieve cost optimization. Also, while expensive, program managers and project leads should explore more parallel and exascale computing technologies. Additionally, mitigating IO limitations needs addressing. IO limitations will grow with increased resolution and ensemble sizes as file systems enlarge along with processing multiple applications simultaneously. Finally, plans should continue towards increasing the R&D to Operations HPC ratio to obtain the proper amount of resources for R&D and T2O activities. This may mean utilizing more cloud vs. on-premises R&D HPC.

Programs and community efforts (e.g. EPIC and UFS) should continue to advance cloud opportunities and leverage community expertise (to include academia and industry) to build software/modeling applications running on cloud platforms. This adds ease of access for non-NOAA users as they would not have to obtain the same restrictive credentials as they would in order to access on-premises machines behind government firewalls. Program and project leads should consider, plan, and budget for operational implementation of some of the applications onto cloud platforms. Some applications, such as RRFS and 3D-RTMA, already have a proof-of-concept and expertise as they have run or are testing applications in the cloud. Similar considerations are underway for WoFS. Longer-range modeling is ideal for surge or spike-computing, such as SFS implementation, especially when running longer length, longer run times, and more computationally-expensive models at reduced cadences. Consideration should also be given to medium-range ensembles such as GEFS. Additionally, application teams can leverage ML and AI advancements to economize cloud compute processes for NWP.

For UFS-R2O, it is absolutely imperative that dedicated compute allocation portfolios are created for each project within RDHPC systems. The ad-hoc approach of receiving "in-kind" contributions from various RDHPC portfolios is not sustainable and hampers the ability for project leads to effectively manage, plan, and execute their allocations. Additionally, streamlining the UFS-R2O allocations by reducing the spreading of allocations across RDHPC machines would save time and decrease the portability risk and time consumption that project engineers and developers face.

Program managers and project leads should build mitigation strategies for utilizing the resources reasonably expected to be available since obtaining the resources needed may not happen. In this case, program managers and project leads should come up with alternative strategies that include leveraging novel computing architectures, methods, and technologies, reducing application scope, adjusting schedules, or a combination thereof. Program managers and project leads should exploit areas such as code optimization, improved parallel scaling, maximize in-core processing, software engineering improvements, and choosing architectures that take advantage of increased GPUs to maximize efficiency. While not a new concept to developers, it is worth entertaining the numerous technical trade-offs or adjustments within application development to lessen computational and storage burdens and facilitate accomplishment of science goals. This may include, for example, reducing cadence, reducing ensemble sizes, adjusting model lengths, compiling code using 32-bit vs. 64-bit for ensembles, performing experiments at reduced resolution, choosing smaller nested grid sizes for regional models, adjusting parameterization, etc. Determining the application's critical path is necessary when weighing technical tradeoffs.

R&R is computational expensive; however, it scales well for surge computing making them good cloud computing candidates. An alternative to the extraordinary expense for GFS/GEFS pre-implementation is to utilize novel methods such as a "replay" approach similar to NASA's Goddard Space Flight Center where they read analyzed fields every 6 hours and nudge the model's background forecasts to the differences between the background and reanalysis. This method would not be as precise, but would only cost approximately 41M core hours compared to the much larger amounts listed in section 2.4.1. Additionally, AI/ML techniques have the ability to address observing system and reanalysis stream discontinuities. Finally, R&R for SFS needs consideration and must be accomplished prior to SFS implementation. SFS could save cost by utilizing the R&R from GFS/GEFS; however, this comes with its own tradeoffs.

Storage is another area that requires serious consideration to utilizing cloud resources such as archival of applications that currently exist on the NOAA Operational Model Archive and Distribution System (NOMADS). Cloud storage, especially leveraging the NODD, is an option, but this must be weighed with procuring more on-premises hardware.

All of this is based on current operational implementation projections with the capabilities documented or planned within the next 5 years. It does not take into account a 10-year approach and more "moon-shot" type estimates of compute necessary to develop the modeling capabilities that NOAA desires to achieve as is explained in the "2017-2018 Roadmap for the Production Suite at NCEP" (A.K.A. "The Roadmap") which takes a more "strategic vision" for NOAA's Environmental Modeling Enterprise. The compute estimates in the Roadmap are much greater than those listed in section 2.4.1. and should be considered.

Program managers and project leads should consider these recommendations in advocating for future compute investments, exploring cost-benefit analysis for cloud technologies, and crafting careful mitigation strategies that weigh critical cost, schedule, and performance tradeoffs. The recommendations presented here are not meant to be authoritative and are not all-inclusive, but rather highlight some key takeaways to help inform planning and decision making to senior leaders in order to obtain the necessary resources and funding to meet research and operational goals. Furthermore, these recommendations can benefit all of NOAA by educating program managers and project leads in requesting the necessary compute/storage resources and funding to meet program/project objectives. Directors should work with higher-level NOAA leaders and the appropriate committees to communicate the needs necessary. This includes working with the NOAA Modeling Team, Modeling Strategy Group, NOAA Observing Systems Council, HPC Board, RDHPC System Allocation Committee, among others, and NOAA/OCIO to develop a consistent cross-Line Office plan. Regardless of the approach and solution decided, all options require extensive planning and multi-year lead-times to advocate and procure the required and substantial resources.

4 References

2017-2018 Roadmap for the Production Suite at NCEP

Advancing Rapid Refresh Forecast System using Cloud-Based High-Performance Computing at NOAA Testbeds

Evolving the NOAA Production Suite on Available Compute Resources (102nd AMS Annual Meeting)

NOAA Cloud Strategy (July 20)

Priorities for Weather Research Report