

National Oceanic and Atmospheric Administration (NOAA) Response to the Science Advisory Board's Environmental Information Services Working Group (EISWG) Statement Concerning the Ongoing NWS Data Dissemination Challenges

[Link to EISWG Statement](#)

1.0 Executive Summary

In June 2021, and with the assistance of their Environmental Information Services Working Group (EISWG), the NOAA Science Advisory Board (SAB) accepted and transmitted to NOAA the, “Statement Concerning the Ongoing NWS Data Dissemination Challenges”. The statement makes four recommendations, with detailed approaches, to inform NOAA's strategy for improving the NWS data dissemination challenges in both the short and long term.

NOAA sincerely thanks the Science Advisory Board and the EISWG for their dedication and contributions to NOAA, as well as their thoughtful and insightful report. NOAA offers responses to each recommendation below, and acknowledges that while plans are in place to address many of these concerns, the implementation of such plans is dependent upon obtaining the appropriate resources.

2.0 Response to Recommendations / Path Forward

Recommendation 1. Design and implement an emergency response: It is imperative the NWS legacy data dissemination systems and associated infrastructure be stabilized, with the immediate goal of providing robust, reliable capacity with backup capabilities that function without the need for data access limits. This will likely require additional critical infrastructure investments that serve primarily as a stopgap (and possibly independent of the Integrated Dissemination Program [IDP]) until longer-term solutions can be finalized and implemented. EISWG encourages NOAA leadership to explore and implement strategies that will lead to increased bandwidth and infrastructure modifications necessary to bring immediate short-term relief to this critical situation.

NWS Response: NWS agrees with this recommendation. In FY 2021 the NWS received an additional \$1.5M to begin addressing the bandwidth issue. The NWS swiftly initiated three procurements to alleviate the restriction. First, the NWS procured an upgraded network card to enable the clustering of firewalls in Boulder to match College Park and enable increased bandwidth throughout. Second, NWS procured upgraded routers as the next step in updating the hardware to support a larger circuit between College Park and Boulder. Third, the NWS procured several load balancers to increase bandwidth by over 50% from 65G to 100G. By the end of Q2 FY22, NWS expects to have the new hardware in place at both IDP Data Centers. Also, by the end of Q3 FY22, the network bandwidth at both IDP Data Centers including the circuits between the two sites will be upgraded to 100G. These upgrades will lessen the current bandwidth constraints for external customers to access IDP services and data. NWS will manage access limits as an appropriate security measure to mitigate against abuse and service attacks. These approaches will be continually adjusted as new methodologies become available, and as resources allow to reduce impacts.

Recommendation 2. Strengthen engagement with the broader Weather Enterprise: The value of NWS public engagement after the PNS was first announced was demonstrated by the

reduced limits implemented in the SCN. NWS should expand the successful engagement with external partners in a public forum to identify immediate, short-term actions all parties can take to provide emergency improvements while the upgrading effort in (1) is underway. The EISWG membership includes representatives from across the Weather Enterprise that are informed of the issues and are engaged in finding solutions, and could co-sponsor or facilitate such a forum.

NWS response: The NWS concurs that direct engagement with the Weather Enterprise has successfully improved our current IDP operations and our plans for the future. Numerous hands-on conversations between NWS leaders, SMEs and individual enterprise members resulted in discovering and removing problematic processes and barriers, to the benefit of NWS operations and all enterprise members. This new level of public-private sector engagement will be continued now and into the future as we work to expand IDP capacity while improving overall service to our external partners and customers. In line with the WRFIA Act of 2017, which states in Sec 401 (3) (B) that the, “EISWG should identify opportunities to improve communications and partnerships among the National Oceanic and Atmospheric Administration and the private and academic sectors,” we appreciate this comment and look forward to further engaging with EISWG in this area. NWS agrees that ongoing engagement with external partners is critical, and appreciates the offer from EISWG to co-sponsor or facilitate forums for this purpose. We further agree that engaging the Weather Enterprise in discussions about improving customer experience are worthwhile. In keeping with our [NWS Partnership Strategy](#), the NWS holds regular Partner Engagement Meetings at least three times per year and routinely communicates with stakeholders at conferences and through emails, news releases, and other informal means. NWS also schedules more subject-specific meetings or webinars as needed such as the June 30, 2021, NWS Partners Webinar focused on Leveraging the Cloud for Numerical Weather Prediction data. We have archived past partner webinars and meetings at: <https://www.weather.gov/wrn/calendar>

Recommendation 3. Prioritize designing and moving to an appropriate scalable architecture: Given the dramatic overall growth in demand, and the intermittent, event-driven surges, it is critical that elements of the architecture be collaboratively designed to address these dynamic data delivery needs. The need for an adaptable, scalable architecture is critical because while future demand is expected to increase, the rate of increase is difficult to predict. The EISWG encourages exploring different options, including:

a. Leverage Content Delivery Networks: A major rapid enhancement for the current data dissemination system may be found in the increased use of Content Delivery Networks (CDN) that can quickly enable greater scalability of the existing NWS data delivery system. Using CDN technology, data files that do not change once produced (bulk model data grids, as one of many examples) can be cached on the edge of the CDN. This removes the need for every user data request to hit the NWS origin servers, providing the potential to greatly reduce the throughput of requests going directly to NWS. Demand will only increase as new models and products are developed, including the just released new version of GFSv16.

NWS response: The NWS recognizes the value of using cloud for data dissemination and has already leveraged Content Delivery Networks (CDN) successfully. The NWS has hosted most of our weather.gov traffic on a CDN which has enabled us to offload roughly 80% of the total internet bandwidth from these websites that would normally be served by NWS data centers. The National Hurricane Center website is also hosted on a CDN, allowing for surges in traffic

during tropical events. In FY21 NWS has undertaken several initiatives that are utilizing the power of cloud computing and the associated CDN capabilities. GIS Viewer and HydroVIS platforms are planning to leverage load balanced and geolocation based caching; and measurable & resilient content delivery through CDN platforms. Selected NOMADS datasets have already been publicly available for dissemination on NOAA's Big Data platform (BDP), i.e. GFS datasets have been available since 2019, HRRR since 2020 with a formal data dissemination project launched in FY21. NWS is working closely with its user-base to ensure that making model data available via the BDP platform and Cloud Service Providers is done in a way that does not disrupt end-users' downstream processes. In parallel to exploring model data delivery via the BDP platform, based on requested funding in FY22, NWS will investigate expanding the use of a CDN provider to serve NOMADS and FTPPRD services to the edge, reducing impacts to IDP on-premise infrastructure. Lastly, NWS is ensuring an end-user/operations support model for any long-term solution delivering model data to the enterprise.

b. Accelerate the migration to commercial cloud networks: EISWG is fully supportive of the IDP, with particular focus on Phase 4, which will require new resources and considerable time. As such, NWS should strongly consider how it could accelerate the migration of their last mile of data dissemination services to secure commercial cloud networks; such networks have demonstrated the ability to facilitate effective and efficient data management, including large-scale distribution. The NOAA Big Data Project, implemented in collaboration with Amazon, Google, and Microsoft web services, provides useful experiences. The last mile delivery of NWS foundational data through commercial cloud environments must also comply with long standing NOAA principles around equal access to all data at no cost and mechanisms to ensure that no entities receive preferential treatment in terms of data available, or speed of availability. In summary, the EISWG endorses the use of commercial cloud networks (identified in the NWS IDP as part of Phase 4) and encourages the NWS to strongly consider reprioritizing such efforts to more rapidly move foundational data to the cloud. This will facilitate increased data distribution and have the greatest impact on reducing data requests from end users to current NWS data systems.

NWS response: The NWS appreciates the SAB recommendation and is committed to following the [NOAA Cloud Strategy](#) that aligns with the federal Cloud Smart strategy, and complying with the principles of free open and equal access to the public. To this end, NWS currently has several initiatives underway that utilize the NOAA Cloud Utility contract in addition to the BDP contract. GIS Viewer and HydroVIS projects are designed to establish a unified GIS platform on a commercial cloud (AWS) that will enable real-time Flood Inundation Mapping (FIM) data to be disseminated. The Damage Assessment Toolkit application that collects and disseminates post weather-event data is already live and operational on an AWS platform. As recommended, NWS is ready to accelerate its cloud migration and has already worked with Forrester to develop cost models for some of the cloud candidates such as NOMADS and MAG as far back as in 2019. However, NWS also recognizes the need to undertake this migration cautiously and deliberately to ensure that security, operational integrity and performance goals are met. NWS also recognized the legacy nature of some of its applications and realizes that a substantial amount of resources must be invested in re-architecting and modernizing. Lift and shift is not a prudent option for many of the legacy applications. We will embrace the SAB suggestions and will ensure a proper emphasis on Phase 4. Our plan is to move in parallel as we transition data access from on-prem to cloud and simultaneously refactor/rearchitect/ready applications for the cloud migration

provided we have the additional resources.

Based on the additional resources proposed in the FY22 President's Budget, the NWS will be able to begin moving to the public cloud in the timeframe referenced in the IDP Plan. However, to move foundational data any more rapidly, the NWS would face a similar trade off decision between increasing the speed of the move to the cloud or leaving some mission-essential applications on the legacy hardware, which is beyond end of life and is not redundant (Phase 3 of the IDP).

Recommendation 4. Enhance user management, product availability announcements, and training programs: Efforts are needed to reduce unnecessarily excessive demands placed on NWS infrastructure by outside entities. The NWS Office of Dissemination should strongly consider the development and distribution of a best practices document for data access. Inclusive to this would be a real-time notification system such as product availability announcements using widely available messaging. Users could then subscribe to receive announcements indicating availability of certain data products as opposed to continually requesting file lists from NWS servers in rapid succession to see if new files are available. To promote compliance with best practices, rather than continuing to allow data users to remain anonymous, it would be beneficial for the NWS to institute enhanced methods that would facilitate fast and definitive identification of individuals, organizations and institutions, to eliminate the challenges caused by their data acquisition methodologies. This capability would allow the NWS, with appropriate public notice to those entities, an opportunity to address and remediate concerns with individuals and organizations that put undue pressure on NWS infrastructure and provide users with ample notice to improve their code in accordance with NWS best practices.

NWS response: The NWS agrees that reducing excessive demand is prudent given the current limitations of IDP. As an example, NWS managed demands during extreme weather events in the past year by distributing guidance on using NWSChat to ensure those who truly need to make use of the system will be able to do so. Similarly we have placed limits on some of the services to detect excessive usage patterns, and we regularly work with users to help them adjust their pattern of usage after being blocked. As for a real time notification system, the NWS is considering the use of a publish/subscribe system that would allow users to choose what information they would like to subscribe to for automatic updates. Additional research is required before the NWS can determine if this is a feasible path forward.

Additionally, the NWS data is used across the globe by other federal agencies, national and international partners, educational institutions, and private enterprises. We agree there are merits to employing a real-time notification system and have been conducting some investigation in that area. We have recently employed SNS (Simple Notification Service) with the NOMADS datasets on the BDP. We are still collecting user feedback on this method of notification and whether it can feasibly be used for the global user community. We do realize that a key ingredient to implementing and maintaining the proactive, real-time user notifications and data-access requires enhanced user and operations support which needs upfront investment. To that end, NWS has requested funding in FY22 and hopes to start implementing it as soon as the funding becomes available. In the mean-time, NWS is implementing a cloud-version of its NWSChat process that allows for user subscriptions to alerts as enabled by the underlying COTS product. In the future, NWS plans to disseminate bulk data on the cloud platform where many of the content distribution and notification

services are available for use without having to spend resources on developing them.

3.0 Conclusion and Acknowledgements

The NWS' mission is to provide weather, water, and climate data, forecasts and warnings for the protection of life and property and enhancements of the national economy. As such, the NWS dissemination systems are crucial to fulfilling that mission. The IDP Plan was drafted with that mission and responsibility in mind, and we are grateful that the SAB review indicates that we have risen to the challenge and appreciate the comments which help focus on the critical aspects of the plan and possible ways to improve as we implement through a continual engagement strategy with our partners throughout the enterprise. NOAA acknowledges and thanks the Environmental Information Services Working Group for their dedication and contributions to NOAA.