

Case Study



High Performance Computing (HPC)
Intel® Xeon® Scalable Processors

Building a Weather-Ready Nation with Intel® HPC in the Cloud

Elastic compute capacity with Intel® Xeon® Scalable processors drives advances in weather modeling and prediction for the National Oceanic and Atmospheric Administration (NOAA)

At a Glance:

- NOAA is using HPC cloud resources powered by Intel® Xeon® Scalable processors to develop an advanced weather model that more accurately predicts severe weather events.
- System integrator GDIT brings together Intel® HPC technologies and Parallel Works software to speed the model's development and provide portability for NOAA's demanding workloads.



GENERAL DYNAMICS
Information Technology



Executive Summary

As one of the world's preeminent weather forecasting organizations, NOAA needs ever-increasing HPC capacity to advance its numerical weather prediction models. NOAA is developing and prototyping its next-generation Rapid Refresh Forecast System (RRFS) on Intel Xeon Scalable processor-based cloud instances at Amazon Web Services (AWS).

Intel HPC solutions are designed to meet the volume, performance, and throughput requirements of workloads such as NOAA's. Intel on AWS allows for granular provisioning of clusters for diverse workloads. Under General Dynamics Information Technology's (GDIT) prime Research and Development HPC System (RDHPCS) contract with NOAA, Parallel Works middleware streamlines management of NOAA's cloud-based resources.

Together, this collaborative effort helps NOAA speed time-to-science while remaining within budget constraints.

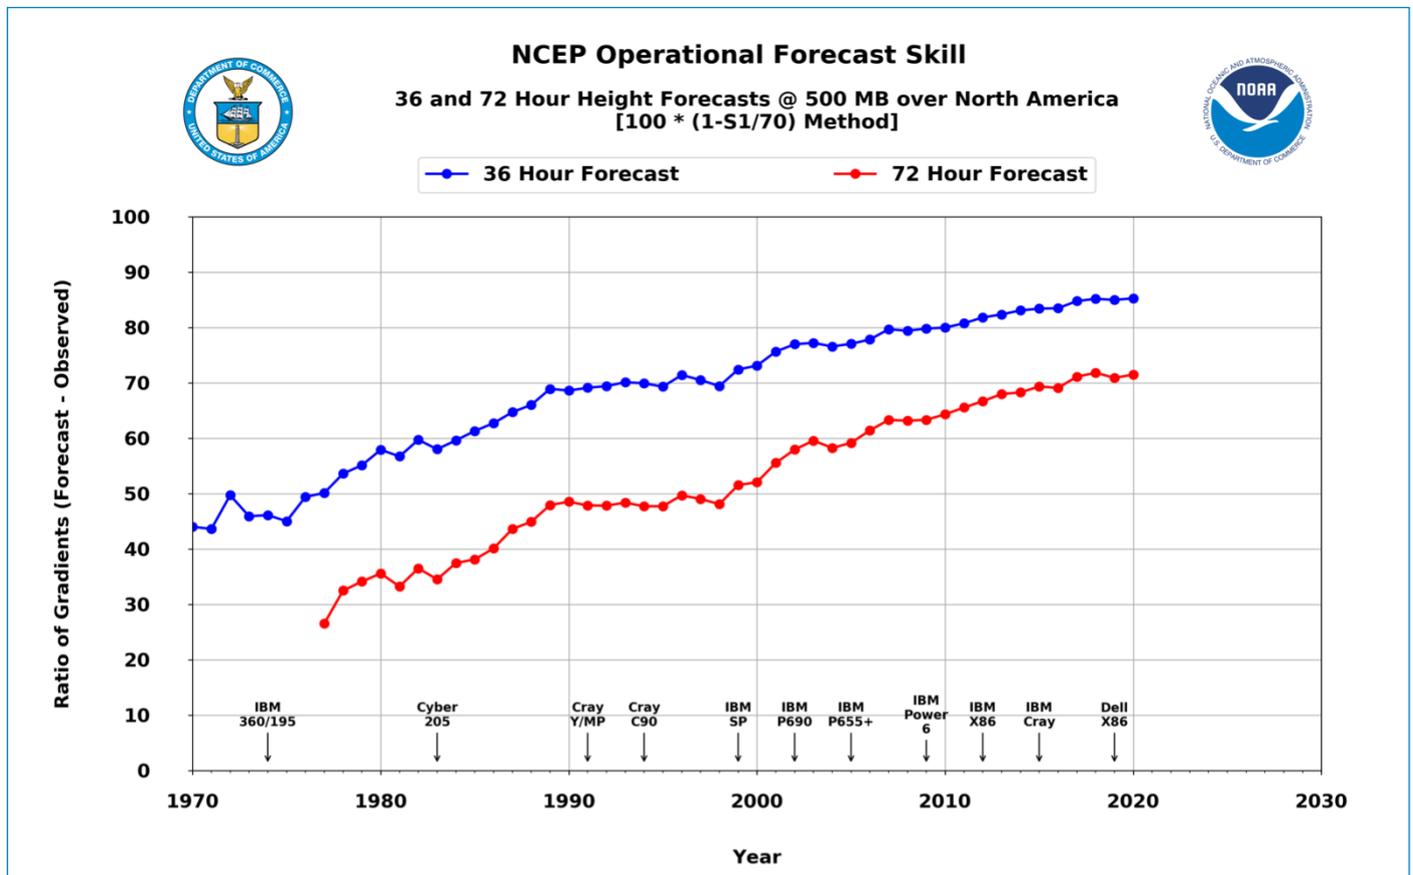
Challenge

From floods to fires, climate change is making extreme weather events more frequent and the consequences more devastating. Climate change also makes timely, accurate forecasts—and the high-performance computers that enable them—more important than ever. High-quality forecasts provide a basis for planning and decision-making, helping save lives and property and enhancing the sustainability of businesses, communities, and regions.

NOAA develops and runs the mathematical weather prediction models that guide U.S. meteorologists in creating their forecasts. NOAA, which includes the National Weather Service, has a mission to make the United States a weather-ready nation—one with the resilience to handle high-impact weather events. NOAA scientists say that increases in HPC capacity are crucial for advancing their weather models and achieving the mission.

HPC for Smarter Forecasts

"HPC and numerical weather prediction have been attached at the hip since the inception of weather modeling over 50 years ago," says Dr. Jacob Carley, a physical scientist at NOAA's Environmental Modeling Center (EMC). "The quality of forecasts is strongly correlated with the rise in HPC performance. Having more HPC resources available allows for us to continually introduce scientific advances, ultimately leading to more timely and accurate forecasts." Figure 1 illustrates this point.



Skill of the NOAA's Global Forecast System over several decades (higher is better). Annotated along the x-axis are various HPC platforms at NOAA. Note the increase in forecast skill with time, facilitated in part by HPC advances that allow for developing and implementing increasingly sophisticated models. Image courtesy of Mallory Row (IM Systems Group and NOAA National Centers for Environmental Prediction (NCEP)/EMC).

Carley is part of a team developing and prototyping the next generation of NOAA's Rapid Refresh Forecast System (RRFS). The RRFS model takes in new weather data and runs multiple times each day, focusing on severe and hazardous weather. It's one of more than 20 models EMC creates and implements.

The RRFS team started with a long list of capabilities it wanted to implement. Many aimed at taking advantage of higher compute performance to improve the model's ability to predict severe weather. Enhancements include expanding the model's geographic coverage and increasing the number of ensemble members (or alternate weather scenarios).

Scientists were also excited by promising algorithmic approaches that would add more complex physics to the weather simulations and increase forecast accuracy for severe and hazardous storm systems.

But NOAA's HPC clusters are in high demand. Developing, prototyping, and testing RRFS would require significantly more resources than developers had available to them through NOAA's on-premises HPC systems.

Solution

Like most research organizations, NOAA is always working to speed innovations through its research-to-operations (R2O) pipeline. To gain the HPC capacity needed for developing RRFS in a timely fashion, NOAA turned to the cloud. While some of NOAA's forecasting models have run in the cloud

before, RRFS would be a test case for building a major model in the cloud.

"There's a growing sense of excitement at NOAA around cloud computing," Carley says. "RRFS is a way to explore cloud capabilities while prototyping NOAA's next-generation, high-resolution system for predicting severe weather."

Intel HPC in the Cloud and On-Premises

"The cloud gives us flexibility to try new algorithms and approaches, which helps us make RRFS more powerful and accelerate its development...We're shaving months or years off the development schedule."

—James Abeles, Computational Scientist, NOAA Affiliate

NOAA is developing RRFS on Amazon Web Services EC2 instances, which offer options to run on 1st, 2nd, and as of Q2 2021, 3rd Generation Intel® Xeon® Scalable processors. AWS has more than 250 EC2 instance types and sizes across its global footprint featuring Intel processors, FPGAs, and other accelerators and technologies. AWS and Intel work closely to deliver offerings tailored to applications such as massively scalable data analytics, artificial intelligence, and the Internet of Things, as well as HPC.

The RRFS team chose Intel processors because its members had extensive experience with Intel architecture and a high

Case Study | Building a Weather-Ready Nation with Intel HPC in the Cloud

degree of confidence in the technologies. They saw value in being consistent with NOAA's in-house infrastructure, much of which runs on Intel platforms.

"We're very familiar with Intel technologies," says Rajendra Panda, a computational scientist at EMC. "We know how to use the Intel Message Passing Interface and the compiler options to get the best performance."

"Having the same architecture on-premises and in the cloud makes it easy to move code back and forth," he adds. "In the early days, we could develop and compile our code on-premises and run it in the cloud. Now, we develop in the cloud."

NOAA uses AWS EC2 instance types based on Intel® Xeon® Platinum processors. Powerful r5.24xlarge instances perform the demanding work of preprocessing. These instances are based on Intel® Xeon® Platinum 8200 processors configured with 768 GB of memory.

Forecasting runs on C5n.18xlarge instances, based on Intel® Xeon® Platinum 8100 processors and 192 GB of memory.

These HPC-oriented AWS instances enable NOAA's codes to take transparent advantage of Intel® Advanced Vector Extensions-512 (Intel® AVX-512) and Intel® Turbo Boost Technology. Intel AVX-512 instructions accelerate performance for a range of scientific, artificial intelligence, and other workloads. Intel Turbo Boost Technology optimizes CPU performance while balancing power and temperature limits.

Performance and Flexibility

"Cloud vendors make it easy to explore new technologies and see how they will benefit our applications. This is a huge benefit in helping us understand what types of technologies we should use in the future and then deploy them quickly."

—Rajendra Panda, Computational Scientist, NOAA Affiliate

Intel on AWS increases solution performance by providing rapid access to the latest Intel HPC innovations.

NOAA's cloud strategy also provides the flexibility to match workloads to the most suitable resources—something that's not always feasible with in-house infrastructure.

"Our on-premises systems tend to be fairly homogeneous, but not every application can use the resources efficiently," says James Abeles, a computational scientist at EMC. "With the cloud, we optimize our cost-performance by allocating different workflow elements to the AWS instance types that run them most efficiently."

Managing Hybrid HPC Resource Portability with Parallel Works

To gain the portability for moving its research workflows between its supercomputers and multiple cloud service providers (CSPs), NOAA deployed Parallel Works hybrid-cloud workflow-management software.

"Parallel Works is a web-based access platform that makes it easy to launch a parallel cluster and submit jobs on any supported cloud," Panda explains. "We use the same

authentication to access the cloud that we do for accessing our in-house systems, which are a great convenience."

GDIT, as the prime systems integrator for NOAA's RDHPC contract, managed the Parallel Works deployment and played an essential role in this joint effort. GDIT is a leading systems integrator for the federal government and has supported the U.S. Department of Commerce NOAA RDHPCS effort since 2010.

GDIT evaluated multiple cloud HPC software frameworks for NOAA and selected Parallel Works to create NOAA's hybrid cloud framework. GDIT teams collaborate closely with NOAA, Intel, AWS, Microsoft, and Google to achieve supercomputing performance levels for NOAA's key weather application codes on all three clouds.

Performance and Budget Management

Parallel Works software also simplifies cost reporting and budget management controls for the RRFS cloud environment.

"We have a lot going on in the cloud," Panda says. "We have multiple compute instances. We store large amounts of data on the cloud, and we move data back and forth between our on-premises and cloud platforms. We also bring publicly available data from the internet onto the cloud each day. Parallel Works' cost-reporting system helps us track our resource usage, which helps us use our resources efficiently and remain within our budget."

The ability to manage its scientists' cloud HPC budgets on a per-project basis was an important requirement for NOAA. The Parallel Works software supports all three major CSPs, and other NOAA teams use the platform to access Intel instances on Google Cloud Platform (GCP) and Microsoft's Azure Cloud.

This transparent multi-cloud capability is key to maintaining a cost-effective, CSP-agnostic environment and making NOAA's vast troves of data available to the public at no cost.

GDIT and Parallel Works developed a baseline HPC configuration optimized for performance using Intel processors on AWS, Azure, and GCP. They based the configuration on a NOAA Finite-Volume Cubed-Sphere, Dynamical Core Global Forecast System (FV3GFS) benchmark application. NOAA scientists calibrated the platform to run scientific models such as FFAIR and the Hazardous Weather Testbed (HWT).

Results

With the elasticity of the AWS cloud and the performance of the Intel Xeon Scalable processors, NOAA has the capacity it needs to significantly expand the RRFS model.

Environmental Modeling Center "Having more computational capacity lets us do higher-resolution forecasts and perform a higher volume of experiments," says Carley. "This allows us to get results and make improvements to the RRFS quickly without sacrificing scientific rigor. It all adds up to being able to make more accurate predictions."

The RRFS is data rich and computationally intensive. "We model the atmosphere in 65 horizontal layers or slabs, and each slab has about 9 million grid cells," says Abeles. "We run

nine members in an ensemble approach to better sample and capture the uncertainties in our forecasts. It's a very large computational problem."

Exciting Science and Faster Research-to-Operations

"With the elastic capacity and high performance of the cloud, we could run and test the prototype model in real time and do exciting, new science at NOAA's flagship testbeds," says Carley. "This collaboration will accelerate our pace of development leading toward a high-quality operational forecast system in the coming years."

—Dr. Jacob Carley, Physical Scientist, NOAA/NWS/NCEP

NOAA's HPC cloud, powered by Intel technologies and Parallel Works software, makes it all possible—and it is improving the project's R2O effort.

"Our HPC cloud lets us scale up elastically to develop and test this system in ways we couldn't otherwise, because the problem size is so large," Abeles says. "The cloud gives us flexibility to try new algorithms and approaches, which helps us make RRFS more powerful and accelerate its development. We can test the system in a more complete way. We can move faster. We're shaving months or years off the development schedule."

The prototype version of the new RRFS model has already been used with several of NOAA's flagship testbed environments and has shown significant results. Forecasters used the prototype in NOAA's 2021 Hazardous Weather Testbed Spring Forecast experiment and FFAIR experiment and characterized both as great successes.

"With the elastic capacity and high performance of the cloud, we could run and test the prototype model in real time and do exciting, new science at NOAA's flagship testbeds," says Carley. "This collaboration will accelerate our pace of development leading toward a high-quality operational forecast system in the coming years."

The new RRFS model is planned for transition to operations at the National Weather Service around late 2023. While the model will initially run on NOAA's in-house supercomputers, its development is providing data and wisdom to inform future decisions about operational forecasting in the cloud.

A Weather-Ready Future

NOAA is looking to move other applications to the cloud, notes Dr. Christina Holt, an atmospheric scientist and software engineer at NOAA's Global Systems Lab (GSL) and NOAA's Cooperative Institute for Research in Environmental Sciences at the University of Colorado. Holt says that GSL is also looking at ways to automate tests of the new weather system.

NOAA is also exploring the use of artificial intelligence to improve weather models and forecasting. Here again, NOAA scientists see value in the cloud.

"NOAA typically keeps its on-premises systems for five years or longer," says Panda. "Cloud vendors make it easy to explore new technologies and see how they can benefit our applications. This is a huge benefit in helping us understand what types of technologies we should use in the future and then deploy them quickly."

Meanwhile, RRFS is an important step toward fulfilling NOAA's vision of weather readiness. "In an environment with more impactful hazardous weather, where communities may be especially vulnerable, it's imperative that the new RRFS provide better forecasts," says Carley. "Our mission is to protect lives and property and enhance the national economy. Early and accurate weather forecast guidance from the RRFS is a significant step toward achieving our mission."

For More Information

- [NOAA](#)
- [Intel HPC](#)
- [Intel Xeon Scalable processors](#)
- [Parallel Works](#)
- [GDIT HPC](#)

Solution Ingredients

- AWS EC2 r5.24xlarge instances for preprocessing
- AWS EC2 c5n.18xlarge instances for forecasting
- Intel Xeon Platinum processors
- Intel Message Passing Interface
- Intel Compilers
- Parallel Works Platform



Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.