

Unlocking the Power of Data, Together

Collaboration between Chief Meteorologist Sunny Wescott, system integrator GAI, AI firm Kamiwaza, and Intel accelerated Wescott's research that can help public and private sector stakeholders better prepare for significant weather events.

At a glance

- Weather researcher and Federal Chief Meteorologist Sunny Wescott wanted to use historical public weather data to inform her barometric pressure research, but the data was in a vast, uninterpretable dataset.
- Through different organizations bringing their strengths and collective vision to the table, scientists and technologists were able to harness information that was out of reach prior to the advent of generative AI (GenAI).
- In less than a month, Wescott had access to the data and GenAI insights that she needed for her research—something that might have taken 5-10 data engineers more than a year to accomplish manually.
- The GenAI solution, which consists of multiple large language models (LLMs), runs on Intel® Gaudi® 3 accelerators and Intel® Xeon® 6 CPUs in the Intel® Tiber® AI Cloud, delivering both excellent inference performance and energy efficiency.

Challenge

Historical meteorological records, obtained from automated sensors around the United States, form the backbone of forecasting models, risk analyses, and early-warning systems. When those records are incomplete or unreliable, emergency managers face shorter warning lead times, less accurate impact predictions, and potentially greater loss of life and property. Where robust historical archives and analytics are in use, outcomes improve dramatically. For example, during Hurricane Harvey, average evacuation times dropped by 40% thanks to plans informed by past flood patterns.¹

Extreme weather trends since 1930 have indicated a major swing in lower lows of barometric pressure and longer-lasting highs, which are causing severe weather at both an increasing frequency and greater intensity. When warning systems and practices are established without relevant historical data, emergency managers may not be able to adequately prepare for weather effects in their location.

Federal Chief Meteorologist Sunny Wescott specializes in national extreme weather hazards and impacts to public and private sector key resources, and cascading threats to critical infrastructure. She is currently working on a research project at the Naval Postgraduate School, focusing on how barometric changes relate to Homeland Security operations. Severe swings in pressure may affect human behavior, which can impact both the frequency of needing to operate, and the staff's wellbeing and ability to operate at full capacity on the days that they're really needed.

Wescott wanted to access the public dataset from Iowa State University to investigate barometric pressure highs and lows across hundreds of specific locations in the U.S., as well as to gain an overall view of how weather has changed over the last 90 years. However, although this dataset contains a treasure trove of site-specific climatic data, it is in an archaic format called GEMPAK, which few contemporary scientists can access.

Solution

To find a way to unpack the GEMPAK dataset, Wescott reached out to Government Acquisitions, Inc. (GAI), a systems integrator she had worked with on previous projects. Emma White, a data scientist at GAI, knew what needed to be done: convert the legacy data format into Parquet, which is a popular, open-source, columnar data storage format optimized for analytical querying and data processing. But a test conversion on a small portion of the dataset told



White she needed more help. She and Wescott connected with Kamiwaza AI, a firm specializing in GenAI solutions. Kamiwaza engineers used AI to scan the GEMPAK file format, gave the AI agent a single background file that described the GEMPAK format, and then started processing the data.

Using the AI agent, the team converted the data from GEMPAK into Parquet—1.3 billion rows and nearly a trillion data points—in just over a week. Cleaning the data, training the AI agent, and generating over 200 graphs took a few more weeks. Comparatively, it would have taken 5-10 data engineers more than a year to accomplish the project manually.²

With the data converted to Parquet, White directed the AI agent developed by Kamiwaza towards the data for analysis. She told the agent to look for specific data, pull the data into a container, and clean it. This involved eliminating null values or barometric values that were clearly inaccurate. The agent understood very well what the team was looking for and took actions to augment the data, such as creating supporting graphs or asking follow-up questions.

“Nobody we knew understood the GEMPAK format. And the dataset was far bigger than any one human would be able to look at and understand. If AI didn’t exist as it does today, I don’t think we would have been able to process it.”

– Luke Norris, CEO, Kamiwaza AI

With such a large dataset and such a lot of analytics, reports, and graphs, the team needed sufficient compute power. The AI agent ran multiple LLM models and needed a system with nearly a terabyte of VRAM. Wescott and her collaborators

turned to the Intel Tiber AI Cloud, where they deployed a system running Intel Xeon 6 CPUs and eight Intel Gaudi accelerators. According to Luke Norris, CEO at Kamiwaza, this configuration enabled the team to “get insights in seconds instead of days or weeks.”³

When optimizing Kamiwaza’s AI engines to run on Intel Gaudi accelerators, Norris credits his company’s membership in [Intel’s Liftoff Program](#) for making the transition quick and easy. The program provided development resources to help rewrite some of Kamiwaza’s code and validate it.

“The AI Liftoff Program provided extreme value to Kamiwaza. It helped us get the resources we needed. I think it was the easiest transition we’ve made.”

– Luke Norris, CEO, Kamiwaza

While inference performance was important, Wescott also valued the energy efficiency of the hardware system and the hosting environment. Intel Gaudi 3 accelerators deliver 40% better power efficiency for inference, compared to the competition.⁴ Also, the Intel Tiber AI Cloud platform integrates advanced cooling technologies, like two-phase liquid cooling, and software-driven sustainability features such as runtime and capacity optimization to help enhance energy efficiency and reduce environmental impact.⁵

Another advantage of working with Intel on the project, according to GAI, is Intel’s robust and highly secure supply chain. “That’s important for both having uninterrupted access to this sort of software and hardware, as well as the security piece, which is important in the public sector,” said White.

Results

Besides bolstering her research paper on barometric pressure variations over time, the platform can help emergency managers and others better understand and prepare for weather events in their location. For example, a new emergency manager might not know whether a forecast for 65 miles per hour winds is normal or record-breaking. Using natural language queries, the manager can ask the AI agent developed by Kamiwaza and GAI to find all of the instances where winds of that nature have occurred at that site and summarize the effects. If the graph shows only three instances and the headlines talk about downed power lines or fallen trees, the manager can prepare accordingly.



“Sustainability was a big part of the conversation [with Kamiwaza]. I wanted to be able to prove to my group that I’m making smart choices using AI.”

– Sunny Wescott,
Federal Chief Meteorologist

Or, if it has happened 2,000 times, the area is likely already hardened against such winds, so the manager can consider it a run-of-the-mill weather event.

The AI platform created by the collaborative team has ramifications far beyond enabling Wescott to access the data she needed for her thesis. It’s an easy-to-use platform that doesn’t require any coding experience to explore data. Building on the key learnings gained from Wescott’s project with the Iowa State University dataset, Kamiwaza is well-positioned to help other public and private sector stakeholders transform massive repositories of real-time and historic data into actionable insights.

For example, the National Data Buoy Center (NDBC) operates the National Buoy System—a huge network of moored and drifting buoys and coastal stations. These buoys serve as critical sources of real-time data for marine weather forecasting, climate monitoring, and emergency response planning. The AI platform developed by Kamiwaza and GAI could mine this data to help improve maritime safety for ships, offshore drilling, and coastal operations.

As another example, utility companies monitor river conditions using a variety of hydrological and environmental sensors to optimize hydropower electricity generation, ensure dam safety, and maintain environmental compliance. This data is stored in a variety of formats, including proprietary binary formats, time-series data, relational databases, text, JSON, and XML. Due to the volume of data and the varying formats, Kamiwaza’s AI agent could significantly accelerate data analysis and enable new insights to drive hydropower operational efficiency.

Solution Ingredients

- [Intel® Gaudi® accelerators](#)
- [Intel® Xeon® 6 processors](#)
- [Kamiwaza AI agent](#)
- [System integration from GAI](#)

*“It’s a really flexible platform. It’s beautiful.
It’s easy to use.”*

– Emma White, Data Scientist, GAI

Learn more by visiting [Intel® Gaudi® AI Accelerator Products and Intel Liftoff Program](#).



¹ FarmstandApp, March 2025, "[7 Ways Historical Weather Data Shapes Smart Future Planning.](#)"

² Interview with Luke Norris, April 1, 2025.

³ Interview with Luke Norris, April 1, 2025.

⁴ Source: Intel Corp., April 9, 2024, "[Intel Unleashes Enterprise AI with Gaudi 3, AI Open Systems Strategy and New Customer Wins.](#)" Power efficiency for Intel® Gaudi® 3 accelerators and Nvidia H100 GPUs was averaged across Llama 7B and 70B parameters, and Falcon 180B parameter models. NV H100 comparison based on <https://nvidia.github.io/TensorRT-LLM/performance.html#h100-gpus-fp8>, March 28, 2024. Reported numbers are per GPU. Vs Intel® Gaudi® 3 projections for LLAMA2-7B, LLAMA2-70B & Falcon 180B. Power efficiency for both Nvidia and Gaudi 3 based on internal estimates. Results may vary.

⁵ Data Centre News, December 2023, "[Vertiv & Intel collaborate on eco-friendly AI cooling solution.](#)"

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.

For workloads and configurations visit www.Intel.com/PerformanceIndex. Results may vary.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. Your costs and results may vary.

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

© 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.

0625/AC/CAT/PDF 365650-001EN