HPCwire Readers' Choice Awards 2025

Welcome!

Below you will find the official 2025 Readers' Choice Award ballot. While there are several categories, the entire voting process should take only a few minutes of your time.

All information is kept confidential and is not shared with any third parties.

Polling for this ballot closes at 11:59 PM PDT on October 3, 2025, followed by winner notifications.

Be sure to subscribe to our weekly newsletter and follow @HPCwire on X for information on the final presentation of these highly anticipated and prestigious awards.

Thank you to all of our *HPCwire* readers for your continued support of our community!

For any questions related to the Readers' Choice Awards, or to report any technical difficulties with voting, please contact shannon@taborcommunications.com.

Please use Chrome to access survey. If you're experiencing technical issues, please clear your browsing history, cookies and cache.

1. Best Use of HPC in Life Sciences

- ◯ **ORCA Computing**, in collaboration with the **Technical University of Denmark (DTU)**, is leading the first-of-its-kind application of quantum-accelerated generative AI to vaccine design, using its PT Series systems to create novel peptide candidates from scratch and demonstrating a powerful new model for life sciences innovation.
- ◯ A **University of Pittsburgh** team has simulated the HIV-1 virus to show how a twist in a critical protein may help it squeeze into a host cell's nuclear pore, causing infection. The research team's simulations

yielded encouraging results, suggesting that simulations can pair with experiments to find new targets for HIV drugs or vaccines.

- **Siam AI**, powered by **DDN**, delivers Thailand's first sovereign AI infrastructure, accelerating national digital transformation by enabling large-scale AI research, healthcare innovation, smart city applications, and localized large language models. Siam AI is driving faster, secure, and scalable AI research and applications across industries.

- **Novo Nordisk**, **Columbia University**, and **Amazon Web Services (AWS)** collaborated to create OpenFold3 with the **OpenFold AI Consortium**. The work required scale-out GPU training on ParallelCluster and Graviton compute for pre-processing.

- Slorado is the first open-source software enabling nanopore sequencing basecalling on **AMD** GPUs, powered by **Pawsey's** Setonix, and developed by researchers from the **School of Computer Science and Engineering, UNSW Sydney**. The work breaks hardware barriers and has helped accelerate genomic research worldwide.

- Researchers at the **New Jersey Institute of Technology (NJIT)** and the **University of Florida** used the Expanse HPC system at the **San Diego Supercomputer Center (SDSC)** to make a groundbreaking discovery about the characterization of hemodynamics in blood vessel sprouts. The work created a foundation for better understanding and prediction of how new blood vessels grow.

- **Scripps Research** partnered with **DDN** to deploy a high-performance, AI-ready data platform to transform cryo-EM workflows. The work enabled Scripps Research to process and analyze cryo-EM data in real time, accelerating protein modeling, vaccine development, and therapeutic discoveries with direct benefits to global health.

- **Novo Nordisk, Columbia University**, and **Amazon Web Services (AWS)** trained OpenFold3 on 256 GPUs, cutting runtime by ~50% and cost by ~55% while scaling to 1M+ alignments/day. The effort accelerates drug discovery by democratizing state-of-the-art protein modeling at a fraction of historical time and cost.

- **The Hormel Institute**'s cryo-EM and cancer-research programs use **VDURA**'s high-performance data platform to process massive imaging datasets in real time, enabling the discovery of a previously unknown molecule implicated in Alzheimer's disease. The work paves the way for clarifying the biological function of Alzheimer's-related molecules and advancing therapeutic research.

- **The University of Wollongong**'s cryo-EM center achieved a significant performance breakthrough by implementing the **VDURA** Data Platform, unlocking 8x more bandwidth to drive its K3 detectors and high-throughput microscopes. The effort allowed the university to stream live data at up to 4.8 GB/s and support more than 120 global researchers via Globus transfers, expanding capacity and accelerating performance to generate atomic-resolution 3D protein structures for drug discovery, disease research, and life sciences innovation.

- A collaboration between **Amazon Web Services (AWS), NVIDIA, AstraZeneca**, and **IonQ** demonstrated the largest quantum-classical auxiliary-field quantum Monte Carlo (QC-AFQMC) simulation to date. The advanced method for predicting chemical reaction pathways achieved up to a 656x speedup and reduced the time to solution for critical drug discovery calculations, offering the HPC

community a compelling example of how quantum acceleration can be integrated into traditional workflows to tackle problems previously beyond reach.

- **Amazon Web Services (AWS)** researchers exploring issues with large protein language model fine-tuning came up with a novel integration of Low-Rank Adaptation (LoRA) and federated learning. The solution enables privacy-preserving AI-driven drug discovery without compromising sensitive data, and has already been adopted by major pharmaceutical companies in the AISB consortium, including AbbVie, AstraZeneca, Johnson & Johnson, and Sanofi.

- **Duke University** researchers introduced a GPU-accelerated adaptive physics refinement framework that enables anatomically realistic, multiscale simulations of cancer cell adhesion at a fraction of the computational cost of explicit models. The work expands the feasible scale of adhesive cancer cell transport simulations, paving the way for new insights into metastasis mechanisms and guiding future therapeutic strategies.

- **Caris Life Sciences** scaled to 2.38M vCPUs on **Amazon Web Services (AWS)** Batch and processed 400,000 RNA-seq samples in 60 hours, proving cloud HPC can deliver national-lab throughput to clinical genomics. The work transforms oncology pipelines by collapsing weeks of data crunching into days, enabling faster precision-medicine insights for patients.

- **Harvard Medical School (HMS)** built SBCloud, an HPCaaS solution, on **Amazon Web Services (AWS)**, to democratize access to HPC resources for structural biologists globally. The work accelerates scientific discovery by providing scale-out HPC with visualization for scientists studying molecular-level interactions for drug discovery.

- **Caris Life Sciences** leveraged **Amazon Web Services (AWS)** Batch and Spot instances to process 400,000 RNA sequencing samples in just 2.5 days. The processing capabilities enable rapid analysis of Caris's comprehensive molecular profiling platform, which combines advanced RNA sequencing with AI algorithms to revolutionize cancer detection, diagnosis, and treatment selection, directly impacting patient care through more timely and precise therapeutic decisions.

2. Best Use of HPC in Physical Sciences

- A multi-institutional team of researchers from **Rice University**, **University of New Mexico**, **University of Utah**, and **University of Texas at Dallas** used **National Science Foundation (NSF)** ACCESS resources on **Texas Advanced Computing Center (TACC)**'s Stampede3 to create the clearest image yet of Yellowstone's magma reservoir, uncovering a volatile-rich cap that helps stabilize the system and redefine eruption risk.

- Using simulations run on the Expanse system at the **San Diego Supercomputer Center (SDSC)**, a study from the **University of Kentucky (UK) Department of Physics and Astronomy** further validated a surprising answer: black holes don't just devour matter, they actively sculpt the very structure of their host galaxies.

- An international team of researchers used data from the new James Webb Space Telescope (JWST) and simulations on supercomputers, including **Pittsburgh Supercomputing Center (PSC)**'s Bridges-2 and **San Diego Supercomputer Center (SDSC)**'s Expanse, to identify and date three brown dwarfs in a globular cluster for the first time.

- Using **Texas Advanced Computing Center (TACC)**'s Lonestar6 supercomputer, Texas researchers run simulations, data analysis, visualization, and machine learning routines to explore the development of rotating detonation rocket engines (RDRE). This new type of rocket engine can improve aircraft, spacecraft, and other transportation systems and is being developed by private and public institutions such as **NASA** and the **Air Force Research Laboratory**.

- By harnessing **NVIDIA** A100 GPUs and OpenACC, researchers from the **Naval Postgraduate School** achieved simulations of rapidly intensifying tropical cyclones more than 10 times faster and with nearly five times more energy efficiency, paving the way for greener next-generation weather prediction.

- Motivated by industrial space exploration’s use of rockets with large numbers of engines, researchers ran the largest compressible Computational Fluid Dynamics (CFD) simulation (> 100 trillion grid points; 10×beyond state-of-the-art) with computational time/energy largely decreased compared with optimized baseline implementation, thanks to the unified addressing on tightly coupled **NVIDIA** CPU–GPU platforms.

- The **SETI Institute** built a GPU backend with **NVIDIA** Holoscan that enables radio astronomers to detect signals of interest in real time using AI, uncover signals previously missed, and reduce false alarms by filtering out false positives.

- Scientists from **Argonne National Laboratory (ANL)**, the **University of Chicago**, the **University of Illinois Urbana-Champaign**, and **Johns Hopkins University** developed RADAR, a high-performance computing framework that enables coordinated, privacy-enhancing multi-messenger follow-up of gravitational wave events in the radio band. This overcame observational constraints and paved the way for efficient analysis in the era of next-generation detectors.

- **Michigan State University (MSU)** conducted a Biophysical Assessment to understand the effects of climate change and mitigation strategies on crop yields and nutrient loss. It used HPC on the **National Center for Atmospheric Sciences (NCAR)**’s Derecho to simulate major crops such as corn, soybean, and wheat under a changing climate in the Community Land Model (CLM).

- Blink is a GPU-accelerated signal processing and imaging pipeline developed by **Pawsey Supercomputing Research Centre** researchers and running on Pawsey’s Setonix supercomputer. Blink enables astronomers to probe for Fast Radio Bursts with the Murchison Widefield Array (MWA) telescope, opening the potential for real-time searches to uncover the origins and physics of Fast Radio Bursts.

- **University of Birmingham** researchers, together with contributors from around the world, can homogenize and assimilate multiple observations into both empirical and dynamical models of the near-Earth space environment in real time, in order to develop new methods of mitigating the impacts of Space Weather on communications, navigation, remote sensing, spacecraft operations, and the sustainable management of near-Earth Space.

- **Lawrence Livermore National Laboratory (LLNL)**, **UC San Diego/Scripps**, and **UT Austin** built a real-time tsunami digital twin using El Capitan that turns ocean pressure sensor data and physics-based models into localized forecasts in under 0.2 seconds, which is about 10 billion times faster than traditional methods. This achievement cuts false alarms and speeds credible warnings.

- Driven by the need for reliable, on-demand HPC resources to handle massive data collection rates and dynamically guide experiments in real time at the recently upgraded **Advanced Photon Source (APS)**, researchers at **Argonne National Laboratory (ANL)** have developed a scalable, resilient, and automated approach for X-ray Photon Correlation Spectroscopy analysis.

3. Best HPC Response to Societal Plight

- **Northwestern University**, **Pittsburgh Supercomputing Center (PSC)**, and the **Texas Advanced Computing Center (TACC)** use PSC's flagship Bridges-2 system to study the full DNA sequences of children up to 20 years old who had suffered sudden death, comparing them to the genes of people who didn't. Using DNA sequences identified by artificial intelligence, the scientists discovered the types of genes and risk-carrying variations that pose a higher probability of sudden death at a younger age.

- **The San Diego Supercomputer Center (SDSC)'s Societal Computing and Innovation Lab (SCIL)** takes a novel approach to creating breakthrough technological innovations to meet complex societal challenges with a commitment to real-world solutions that leverage integrated workflows, next-generation data and AI, cutting-edge science, and advanced digital infrastructure. A defining feature of SCIL is its Immersion Studio, which leverages the AI-readiness of scientific data to power new modes of teaching, training, decision-making, and public engagement.

- **The Met Office** leverages Azure HPC to deliver critical weather insights during emergencies, supporting public safety and resilience. Key technologies include **Microsoft Azure** supercomputing, **HPE** Cray EX supercomputers, **AMD** EPYC processors, InfiniBand networking, and high-performance active data archives. The work supports real-time modeling during extreme weather events that help communities prepare and respond to weather emergencies effectively.

- **California Polytechnic State University** used the **National Center for Supercomputing Applications (NCSA)** Delta GPUs to create an AI assistant that utilizes machine learning and was trained on a database of missing person data, using more than 1,000 GPU hours to train their diffusion models. The AI application can be used in the field during search and rescue missions to help mitigate this plight in the United States, where approximately 600,000 people are reported missing each year, and a substantial number are not found in time to save their lives.

- **The National Science Foundation (NSF)**'s NHERI DesignSafe uses HPC to enable scientists and engineers to manage data that helps society become more resilient to life and property-threatening natural hazards. DesignSafe's computational resources, software, and expertise have empowered 10,000 users with tools for managing, analyzing, and sharing data on an expanding list of natural hazards, including earthquakes, hurricanes, inland flooding, tornadoes, landslides, and wildfires over the last decade.

- **Lawrence Livermore National Laboratory (LLNL)**, **UC San Diego/Scripps**, and **UT Austin** used El Capitan to build a real-time tsunami digital twin that turns ocean

pressure sensor data and physics-based models into localized forecasts in under 0.2 seconds, which is about 10 billion times faster than traditional methods. The output cuts false alarms and speeds credible warnings for coastlines exposed to earthquakes on the Cascadia Subduction Zone, including Washington, Oregon, and Northern California.

- The HENS (Huge Ensembles) model and forecasting process is able to capture the long tails of climate distributions, which is something traditional ensemble models are unable to do. The work requires the orchestration and automation of complex data management tasks using Globus by the **Lawrence Berkeley National Lab (LBNL)**, with help from **NVIDIA**.
- The Hurricane Analysis and Forecast System (HAFS) represents a critical modernization initiative for the **National Weather Service's (NWS)** hurricane forecasting capabilities, requiring significant cloud computing resources. By working with the **National Oceanic and Atmospheric Administration (NOAA)** to accelerate HAFS on **Amazon Web Services (AWS)**, this effort enables more accurate and timely hurricane forecasts that help protect lives, property, and communities during increasingly active storm seasons.
- **NVIDIA**'s Earth 2 allows decision-makers to mitigate the effects of climate change by democratizing the training and running of AI weather and climate models and interacting with the predictions through rich digital twins. Earth 2 is already being used to better quantify the risk of extreme weather through generating ensembles of future weather at unprecedented scale and spatial resolution.
- **BurkinaBioinfo (BBi)** established the first dedicated bioinformatics HPC platform in Burkina Faso, enabling local and regional scientists to process complex genomic data to address pressing challenges in health, agriculture, and biodiversity. BBi has trained over 280 scientists from seven West African countries, enhanced regional capacity to respond to disease outbreaks, improved crop resilience research, and reduced dependency on external facilities for data analysis.

4. Best Use of HPC in Energy

- Researchers at the **National Center for Supercomputing Applications (NCSA)** and the **University of Illinois Urbana-Champaign** developed a deep learning operator, based on a virtual sensing digital twin and trained on the **NVIDIA** GH200-powered DeltaAI HPC cluster, to monitor inaccessible nuclear reactor locations in real time. The work delivers predictions of critical and previously unmeasurable parameters 1,400× faster than traditional simulations, surpassing the limits of conventional sensors and AI methods.
- With help from **Argonne National Laboratory (ANL)** and **Lawrence Berkeley National Laboratory (LBNL)**, the **DIII-D National Fusion Facility** team, led by **General Atomics** developers, has implemented new workflows using **Globus** platform services to streamline and automate workflows. The work enables researchers to perform higher fidelity reconstructions and near-real-time particle tracking within the DIII-D research program.
- **Energy Aware Runtime (EAR)** enables more than a dozen HPC and AI data centers to achieve up to 50% energy savings while increasing performance, delivering measurable impact on both operational

costs and sustainability goals. EAR is a vendor-agnostic, runtime software that combines full-stack energy monitoring, real-time optimization, and intelligent workload management for HPC and AI workloads.

- Two graduate students from **Tennessee State University** leveraged **National Science Foundation (NSF) ACCESS** resources on Jetstream2 to reveal how biochar and optimized nitrogen use can improve soil health, reduce emissions, and support sustainable biofuel production. This research points toward scalable agricultural practices that can cut greenhouse gas emissions, reduce farmer reliance on costly fertilizers, and strengthen bioenergy and food security for society at large.

- **Shell** has successfully leveraged a cloud-based plus burst HPC capability to gain access to additional HPC capabilities and expanded their HPC capacity by 3x to 5x, with plans to scale to 10x+. The hybrid HPC approach combines on-premises infrastructure with **Amazon Web Services (AWS)** GPU cloud bursting capabilities for seismic processing, emerging AI, and other innovation use cases. This extra capacity minimizes project delays, enables R&D efforts to deploy new algorithms and workflows, and accelerates the timeline for decision making.

- **Penn State** researchers used **Pittsburgh Supercomputing Center (PSC)**'s Bridges-2 supercomputer to add a time dimension to seismic measurements, as well as to analyze how oil damps down the loudness of sound traveling through it. Their preliminary analysis suggests that hidden rock structures in oil reserves commonly prevent all the oil from being pumped out.

- **Xcimer Energy's** IrisNet is a deep learning model trained on extensive HPC-generated simulations, enabling rapid optimization of laser beamlines for inertial confinement fusion. The effort relies on Iris, a novel optical simulation code built on the MFEM finite element library, executed in parallel on **Amazon Web Services (AWS)** Graviton4 clusters deployed with Amazon FSx for Lustre, orchestrated by AWS Parallelcluster.

- A plant-validated, industrial-scale transient multi-physics CFD model developed by researchers from **École de Technologie Supérieure (ÉTS)** on **Compute Canada** clusters now guides operating set points, loading strategies, and heating-element layout design for electric furnaces, accelerating the transition from gas-fired units to renewable electricity. The work provides actionable, replicable guidance that enhances thermal uniformity and throughput while reducing fossil-fuel dependence and CO_2 emissions, leveraging Québec's abundant renewable electricity.

- **BP** is transforming its exploration capabilities by upgrading its HPC Center in Houston. This strategic enhancement is focused on improving time to insights, enabling more predictive exploration, and providing BP with the ability to run multiple scenarios from the start of the seismic value chain. With GPU-based computing, BP can perform Elastic Full Waveform Inversion up to seven times faster, dramatically improving both the speed and quality of seismic imaging.

- Using **Amazon Web Services (AWS)** Graviton processors across hundreds of EC2 Spot Instances, S-Cube and Petrobras processed 3.3 TB of seismic data (700 km² area, 10 km depth, 50/25m grids) using TTI acoustic and elastic wave equations (the latter being 5-10x more demanding). The researchers were able to complete each FWI iteration in 72 minutes while processing 325 shot gathers, demonstrating the cost-effective, high-fidelity subsurface imaging possible for large-scale cloud-based FWI workflows.

5. Best Use Of AI Methods for Augmenting HPC Applications

- A team led by researchers from the **University of Maryland** has used two supercomputers, **Pittsburgh Supercomputing Center (PSC)**'s Bridges-2 and **Johns Hopkins University**'s Rockfish, to create and run DAWN, an AI-enhanced interactive program that allows farmers to explore different crop management options given the weather expected from national climate predictions. DAWN promises to be a user-friendly tool that can be employed to minimize the risks associated with planning crops.
- Researchers from the **National Center for Supercomputing Applications (NCSA)** and the **University of Illinois** have developed a video diffusion Generative AI workflow, trained on DeltaAI, to design innovative meta-materials made of multiple material components, a challenge that classical computational design methods currently deem unsolvable. This project is set to revolutionize functional material design, enabling the creation of customized materials with unprecedented mechanical properties and behaviors, and driving innovation across industries.
- **Amazon Web Services (AWS)** researchers achieved 10,000x speedups over FEM in predicting the response of materials (so-called ball-impact elastodynamics tests and simulations) by fusing ML surrogates with physics-informed training. The work enables aerospace and automotive engineers to evaluate safety-critical designs in real-time, interactively rather than waiting hours or days.
- **PhysicsX** has developed revolutionary Large and Deep Physics Models (LPMs/DPMs) that accelerate computer-aided engineering (CAE) simulations by 10,000x while improving accuracy, enabling engineers to explore hundreds of thousands of design iterations per week. The technology is integrated into an end-to-end platform called Flux and has demonstrated a transformative impact across multiple industries, including reducing manufacturing defects by 70%, improving operational downtime by 20%, and more.
- **University of Birmingham** researchers have developed a high-throughput biodiversity screening kit, now being adapted for in-field use, that enables real-time, high-resolution monitoring of community biodiversity. The effort leverages the University's BlueBEAR cluster (**Lenovo**, with **Intel** CPUs and **IBM** Storage Scale), including **NVIDIA** A100 GPUs and HPC, to process terabytes of biodiversity and environmental data.
- **Pennsylvania State University** researchers and others have developed a deep learning-assisted microstructure reconstruction framework to study the hydration of cement in space. Using this model, researchers were able to generate high-fidelity virtual samples of cementitious samples cured aboard the ISS, which paved the way for creating virtual testing of space cement.
- **Duke University**'s HarVI integrates machine learning with high-performance computing to deliver real-time, patient-specific coronary intervention planning for coronary artery disease, reducing turnaround from days to under 90 minutes and transforming clinical decision support. HarVI combines one-shot training of reduced-order CFD simulations on HPC clusters with a machine learning backend and immersive XR front-end.
- **Toyota Research Institute (TRI)** adopted the newly integrated **Amazon Web Services (AWS)** Batch with SageMaker Training Jobs to orchestrate priority-queued, fair-share, GPU-accelerated training for its large behavior models. The solution dramatically boosts utilization, reduces idle time, and lets scientists focus on model innovation rather than infrastructure.

- **The European Centre for Medium-Range Weather Forecasts (ECMWF)** has developed portable, scalable workflows that enable training and inference at scale on EuroHPC and other European systems, with rapid deployment across architectures to accelerate AI development for weather, climate, and beyond. The work will support the next generation of AI weather and climate solutions, harnessing the potential of European HPC, and in particular, the upcoming AI factories, to support climate adaptation and resilience.

- **Talos Innovation**, assisted by ***n*AG**, have developed Talos, a machine-learning computational enhancement layer that augments existing physics solvers to significantly boost simulation speed without altering core workflows. TalosAps easily integrates into existing IP. It improves the economics of simulation by leveraging GPUs and TPUs to deliver significantly faster results, while providing an increase in compute performance per dollar compared with standard HPC setups.

- Researchers at **Lawrence Berkeley National Laboratory (LBNL**) are using HPC from **HPE** to train AI models to analyze patient records at scale in order to speed up the detection of warning signs and health conditions. Specifically, they are analyzing structured and unstructured clinical data to identify key risk factors for conditions that are often missed by traditional methods, such as suicide and obstructive sleep apnea.

- Amazon SageMaker is a fully managed machine learning (ML) service by **Amazon Web Services (AWS)** that provides a comprehensive set of tools (e.g., AWS Batch, Amazon EC2 GPU, Amazon EFS, Amazon S3) to build, train, and deploy ML models and generative AI applications. Using the tools, advanced ML models, such as U-Nets and Fourier Neural Operators (FNOs), can deliver high-fidelity predictions at speeds around 10,000 times faster than FEM simulations.

6. Best Use of HPC in Financial Services

- Jump Trading leverages **DDN**'s 400NVX2 QLC systems to empower quantitative analysts to operate independently. By decreasing processing latency by 10x and enabling instant results, Jump Trading has dramatically increased quant productivity and innovation, giving the firm a decisive competitive advantage in high-frequency trading.

- **Cornell University** leveraged **National Science Foundation (NSF) ACCESS** allocations on **Texas Advanced Computing Center (TACC)'s** Stampede3 to model how debt relief design and mortgage firm behavior shaped the U.S. foreclosure crisis. This produced insights that earned international recognition in regulatory economics.

- **Qube Research & Technologies (QRT)** needed to extract a complicated business from a large investment bank and turn it into an independent firm in 100 days, which QRT did using **Amazon Web Services (AWS)** migration capabilities. Specifically, at setup, QRT migrated from an on-premises data store to a Greenfield AWS S3 store, leveraging AWS Snowball devices to migrate data, and using POSIX filesystems in the form of AWS EFS and FSx for backwards compatibility.

- Built by **LDA Technologies** on their newest Neo X platform, Neo Tap X can contain up to three of the latest **AMD** Versal Premium HBM FPGAs, and utilizes the latest AMD CPUs in a compact 1RU-tall HPC server configuration, delivering industry-standard precision with 100-picosecond time stamping

accuracy and more, to meet the growing high-bandwidth demands of large financial institutions and trading firms.

- **Amazon Web Services (AWS)** researchers initiated agent-based market (ABM) backtesting on AWS. An ABM for equity market backtesting and simulation can offer insightful observations into market dynamics and the effectiveness of trading strategies under various conditions. In this case, by scaling ABM simulations on AWS, researchers overcame data bias and created robust synthetic markets for strategy backtesting.

- **Balyasny Asset Management (BAM)** sought to empower its approximately 160 investment teams to conduct cutting-edge research across six strategies. The company harnessed **Amazon Web Services (AWS)** Batch to supercharge their research capabilities, enabling teams to efficiently tackle highly parallelized workloads. The work allowed them to surge throughput and research velocity, becoming the blueprint for financial services HPC in the cloud.

7. Best Use of HPC in Industry (Automotive, Aerospace, Manufacturing, Chemical, etc.)

- **Škoda Auto** is advancing its vehicle design process by continuously improving the integration of Computer-Aided Engineering (CAE) simulations and HPC to significantly reduce development time and costs, minimize negative environmental impact, and achieve more efficient and precise designs that boost vehicle performance and energy efficiency. The work uses an **HPE** Cray Supercomputing EX4000 with **AMD** EPYC Genoa CPUs, along with other HPE technologies.

- Researchers at **Arizona State University** and the **University of Illinois Urbana-Champaign** developed a high-fidelity 3D modeling framework on **National Science Foundation (NSF)** ACCESS resources to predict and mitigate reflective cracking in airfield pavements in order to improve aviation safety and infrastructure resilience. This work will help the U.S. Federal Aviation Administration (FAA) and engineers design longer-lasting, safer airfield pavements, reducing maintenance costs and more.

- **Toyo Tire** is using HPC and AI to speed up design and cut simulation time by up to 50%, resulting in enhanced EV tire longevity and efficiency. The work makes use of technologies including **HPE** GreenLake Hardware, HPE Cray XD2000, HPE Apollo 2000 Gen10 Plus System, HPE Apollo 4200 Server, HPE ProLiant DL360 server, HPE Parallel File System Storage, and **Mellanox** HDR InfiniBand.

- Researchers from the **University of Illinois**, the **Massachusetts Institute of Technology (MIT)**, **Sandia National Laboratories (SNL)**, and **National Center for Supercomputing Applications (NCSA)** have developed a novel generative AI model, trained on NCSA's Delta system. The model enables inverse design of complex patterned polymers by rapidly generating multiple high-fidelity manufacturing solutions from desired pattern images, and advances AI-driven materials design through HPC.

- **Schaeffler Technologies AG** slashed software deployment time from two days to 45 minutes, and Virtual Engineering Workbench (VEW) setup time was reduced from 30 minutes to six minutes, powering hundreds of cloud-hosted test benches on **Amazon Web Services (AWS)**. The effort accelerates software-defined vehicle development, shortening feedback loops and enabling safer, smarter mobility systems.

- **École de Technologie Supérieure (ÉTS)**, in collaboration with **SAFRAN,** used HPC to run detailed simulations of jet plume contrail formation. The work combines a new microphysical model with an in-house high-resolution Computational Fluid Dynamics (CFD) code to support aerospace climate impact studies.

- **Microsoft** and **Quantinuum** created 12 logical qubits and demonstrated a hybrid end-to-end chemistry simulation, combining HPC and quantum computing. This effort is leading to breakthroughs in quantum chemistry simulations, paving the way for new materials and industrial processes.

- **University of Illinois, Pittsburgh Supercomputing Center (PSC)**, and **National Center for Supercomputing Applications (NCSA)** researchers used NCSA's Delta supercomputer to reconstruct atomic images from electron microscope data, and PSC's Bridges-2 to carry out molecular dynamics simulations that identified scientifically interesting phenomena likely to occur. The observed phenomena affected the thermal vibrational behavior of the atoms, pointing the way toward designing new, more heat-resistant and thus smaller electronics.

- The **NASA Jet Propulsion Laboratory** selected **Microchip**'s PIC64-HPSC for its High-Performance Spaceflight Computing (HPSC) platform. The platform delivers deterministic and fault-tolerant HPC, onboard AI, advanced sensor fusion, and autonomous decision-making to aerospace missions.

- **Merck** is leveraging its new HPC resources to transform drug discovery processes, helping to identify new drug candidates more effectively and to ultimately bring new treatments to patients more quickly. A new central HPC solution provides computing resources to support and accelerate its teams' research activities, using AI to analyze vast datasets and enabling researchers to identify potential drug candidates more quickly and accurately.

- **Schaeffler Technologies AG** leveraged Azure Storage and HPC for discrete manufacturing, enabling scalable simulations and data management. The work enhanced manufacturing precision and operational agility.

8. Best Use of High Performance Data Analytics & Artificial Intelligence

- Researchers from the **Ethiopian Artificial Intelligence Institute** harnessed the power of HPC and advanced AI to address two critical challenges in Ethiopia: drought risk prediction and Land Use/Land Cover (LULC) classification. Leveraging vast volumes of multi-source climate, hydrological, and satellite imagery data, AI models provide highly accurate, near real-time drought risk forecasts and detailed land use classifications. The work enables rapid data processing, fine-grained spatial mapping, and predictive insights that are vital for disaster preparedness, agricultural planning, and sustainable land management.

- **Harvard** researchers have used large-scale LLMs to analyze 2.2 billion daily geotagged tweets since 2010, creating the most extensive dataset to date of 46 dimensions of Human Flourishing across U.S. states and counties. The project produces high-resolution human flourishing measures, open-source LLMs, interactive dashboards, and analytical workflows that offer granular well-being insights to policymakers, organizations, and the general public, enabling data-driven interventions and community empowerment to address disparities in well-being and improve quality of life.

- **Argonne National Laboratory (ANL)**, **Saudi Aramco-Dhahran**, and **Aramco Americas-Detroit** developed a generative deep learning and artificial intelligence framework for accelerated inverse molecular design of novel high-performance transportation fuels. The proposed data-driven framework streamlines innovative fuel design processes and addresses the limitations of traditional methods by harnessing the power of generative AI and HPC.

- **Riskthinking.AI** built a climate risk digital twin on **Amazon Web Services (AWS)** that runs hundreds of billions to trillions of simulations across petabytes of data for financial institutions. It equips banks, utilities, and regulators with defensible climate risk metrics, shaping smarter investment and resilience strategies.

- A team from the **University of Michigan** used **Pittsburgh Supercomputing Center (PSC)**'s NSF-funded Neocortex supercomputer to simulate hypermutators, which are mutants with the ability to mutate more quickly. Neocortex allowed the scientists to expand their artificial populations from the tens of thousands in earlier simulations to billions, helping to explain the relative lack of hypermutators in real organisms.

- **University of Birmingham** researchers leveraged multi-objective Bayesian optimization and GPU-accelerated (OpenMP offload with MPI) high-fidelity fluid-structure simulations to unravel and optimize the vortex-dominated unsteady wake dynamics of butterfly swarm flight. The work showcases how cutting-edge AI and high-performance data analytics can drive breakthroughs in bio-inspired flight design and control.

- The Community Research Earth Digital Intelligence Twin (CREDIT) framework, developed at the **National Science Foundation (NSF) National Center for Atmospheric Research (NCAR)**, offers a flexible, scalable, foundational research platform for training and deploying AI weather models, providing an end-to-end pipeline for data preprocessing, model training, and evaluation.

- Researchers at the **European Institute of Oncology** and the **Monzino Cardiology Center** are creating predictive, prognostic, and diagnostic computational models based on the interactions of protein structures. The project is expected to improve the development of therapies and treatments through the identification of potential molecular candidates and interaction mechanisms using AI, leading to greater efficiency in healthcare activities.

- A multi-institutional team from the **University of Montana** and the **University of Arizona** used **National Science Foundation (NSF)** ACCESS resources on Jetstream2 to launch MDRepo, the first open-access repository for molecular dynamics simulations, creating a transformative platform for protein research, drug discovery, and AI model training. By centralizing and enabling large-scale analysis of biomolecular simulations, MDRepo accelerates drug discovery and advances biomedical research.

- **Santa Clara University**'s Computer Science and Civil Engineering departments deployed **VDURA**'s V5000 all-flash platform with GPU-dense compute to power real-time traffic incident captioning and AI-driven stream-flow prediction. The work advances public safety, environmental stewardship, and academic innovation at an enterprise scale.

- Scientists from **Argonne National Laboratory (ANL)**, the **University of Chicago**, the **University of Illinois Urbana-Champaign**, and the **University of British Columbia** developed a hybrid AI framework that combines Physics-Informed Neural Operators with generative diffusion models to

accurately simulate incompressible and resistive magnetohydrodynamic turbulence up to unprecedented Reynolds numbers. The work is critical in fields including fluid dynamics, astrophysics, and plasma physics research.

- ◯ Privacy-protected Automatic Speech Recognition systems for children allow researchers to identify speech and language problems in young children and enable faster interventions to help them. The research, which uses **Texas Advanced Computing Center (TACC)**'s Lonestar6 and Corral supercomputers, will help enable a safer digital future for children.

9. Best HPC Storage Product or Technology

- ◯ Phison Pascari X-Series Enterprise SSDs
- ◯ HPE Cray Supercomputing Storage Systems E2000
- ◯ Google Cloud Managed Lustre
- ◯ VAST DataStore
- ◯ Amazon FSx for Lustre
- ◯ VDURA Data Platform V5000
- ◯ Post-Quantum Cryptographic (PQC) Technology
- ◯ Mediaflux for AI
- ◯ BeeGFS 8
- ◯ Starfish Storage
- ◯ Swiss Vaults "VaultFS"
- ◯ xiRAID

10. Best HPC Cloud Platform

- ◯ Altair HPCWorks
- ◯ Rescale
- ◯ Boost Run Platform
- ◯ Massachusetts Open Cloud (MOC)
- ◯ Parallel Works ACTIVATE AI
- ◯ AWS Parallel Computing Service (PCS)
- ◯ Commercial Cloud Resources Provided via CloudBank 2
- ◯ Amazon Web Services

- Google Cloud
- Parallel Works ACTIVATE High Security Platform

11. Best AI Product or Technology

- Amazon EC2 Trn2 UltraServers (Trainium2)
- Amazon Nova
- ANEMOI Open-Source ML Framework for Weather and Climate
- Atlas AI Cluster
- Cerebras Systems Inference
- DDN Infinia
- Lenovo ThinkSystem SR780a V3
- NVIDIA GB300 NVL72
- NVIDIA Warp
- Penguin Solutions OriginAI AI Infrastructure Solution
- VAST AI OS

12. Best Use of HPC in the Cloud (Use Case)

- **The National Science Foundation (NSF)**-funded Early-concept Grants for Exploratory Research (EAGER) project at the **San Diego Supercomputer Center (SDSC)** provides in-depth support to research groups. The effort concentrates on using the **NVIDIA** DGX cloud platform with a focus on optimization of system setups, performance monitoring, and determining the best ways to run National Artificial Intelligence Research Resource (NAIRR) pilot projects on the resources.
- **Arm** has implemented cloud-based Electronic Design Automation (EDA) enabled by HPC. By leveraging **Amazon Web Services (AWS)** and **Exostellar**'s Infrastructure Optimizer, Arm reduced backend EDA workload costs by 40% while maintaining workload integrity. The solution enables 65% of backend workloads' job runtime on spot Instances, which traditionally required on-demand instances due to their stateful, long-running nature.
- **The National Oceanic and Atmospheric Administration (NOAA)** used **Parallel Works** ACTIVATE to move complex, time-sensitive HPC workloads from traditional on-premises systems to **Amazon Web Services (AWS)** cloud infrastructure, while maintaining the extreme reliability and performance standards required for life-saving hurricane forecasting operations. The success of this effort demonstrates that mission-critical weather forecasting can reliably transition from traditional HPC systems to cloud infrastructure.

- **Amazon Web Services (AWS)** researchers were able to accelerate their engineering design and simulation workloads by integrating **Ansys** Gateway with AWS ParallelCluster. The effort enabled them to auto-scale CAE clusters, slashing idle costs and boosting throughput with zero re-platforming. That, in turn, allows engineering teams to run more and larger studies faster, accelerating design cycles and innovation without overprovisioning.

- **Marvell** has embraced a cloud-first approach to EDA by collaborating with **Amazon Web Services (AWS)** for electronic design automation in the cloud. This migration enables Marvell to accelerate silicon design using the advanced and scalable compute capabilities of AWS. The bursty nature of advanced node silicon design workloads is optimally addressed by the virtually unlimited scale and elasticity enabled by AWS compute infrastructure.

- **Dun & Bradstreet** re-platformed their legacy mainframe and Netezza systems to Ocient on **Google Cloud**, achieving a 96.3% reduction in processing time and a projected cost savings of $3.4 million, while modernizing a computationally intensive data processing environment.

13. Best HPC Programming Tool or Technology

- Globus
- Julia
- NextSilicon Maverick-2
- Spack 1.0
- Qbit Bridge – Hybrid Quantum-Classical Workflow
- XGBoost Accelerated by NVIDIA Grace Hopper
- NVIDIA Warp
- Chapel 2.5

14. Best HPC Server Product or Technology

- Micron 9650 NVMe SSD
- Amazon EC2 P6e-GB200 UltraServers (Blackwell)
- HPE Cray Supercomputing EX4000
- HPE ProLiant Compute XD685
- GB200 NVL4
- Q.ANT Native Processing Server
- GigaIO | SourceCode Gryf

- GigaIO SuperNODE
- Lenovo ThinkSystem SC750 V4
- NextSilicon Maverick-2 Accelerator

15. Best HPC Interconnect Product or Technology

- CN5000 Omni-Path Product Family
- NVIDIA Quantum-X Photonics
- NVIDIA Quantum-X800 InfiniBand Networking Platform
- Elastic Fabric Adapter (EFA) v4
- AWS Nitro System
- GigaIO FabreX
- HPE Slingshot Interconnect

16. Best HPC Collaboration (Academia/Government/Industry)

- The tri-party team of **Petrobras**, **Universidade Federal Fluminense**, and **Amazon Web Services (AWS)** engineered a Spot-savvy hybrid framework that cut costs up to 91% while preserving performance for large reservoir simulations. The collaboration eliminates queue bottlenecks and expands compute access to scientists during peak demand without additional CapEx.
- A collaboration between the **Louisiana Coastal Protection and Restoration Authority** and the **Pittsburgh Supercomputing Center (PSC)** combines high-performance computing, user-friendly visualization, and more. The new tool is available for government officials and academic researchers to investigate the likely effects of various environmental policies and projects, helping to avoid coastal degradation and maintain the resources needed for the area's economy to thrive.
- This collaboration between **Parallel Works**, the **Defense Innovation Unit (DIU)**, and **High-Performance Computing Modernization Program (HPCMP)** transforms how the Department of Defense accesses secure, scalable computing resources. The collaboration reduces secure environment deployment time from 18+ months to days while saving each program over $1 million in accreditation costs.
- **The High-Performance Software Foundation (HPSF)** brings together collaborators from industry, academia, and government to help sustain performant, portable HPC software projects. HPSF sponsors activities that enable collaboration among members and projects. In its inaugural year, it established a formal governance structure and spawned working groups around HPC-focused continuous integration (CI) and benchmarking.
- **ORCA Computing** collaborated with **NVIDIA**, **Poznan Supercomputing** and **Networking Center (PCSS)**, and **Imperial College London** to demonstrate a distributed quantum neural network in a

hybrid HPC data center environment. The effort is a first of its kind to integrate multiple quantum processing units (QPUs) and GPUs for multi-user access.

- A multi-disciplinary team of researchers from the **University of Illinois at Urbana-Champaign**, the **Mayo Clinic**, **Eleuthra Photonics**, **Carle Health**, and **the National Center for Supercomputing Applications (NCSA)**, with the funding of an ARPA-H contract from the U.S. Department of Health and Human Services, is using HPC resources to improve patient outcomes by integrating new optical imaging technologies and AI into the surgical process of removing cancerous tumors.

- By integrating resources from **Argonne Leadership Computing Facility (ALCF)**, **Oak Ridge Leadership Computing Facility (OLCF)**, and the **National Energy Research Scientific Computing Center (NERSC)**, OpenCosmo is designed to help researchers efficiently access and process the extensive datasets needed for cosmological studies. It is built on Globus platform services, using Globus Flows to manage data transfers and analyses across HPC sites.

- A team of **University of California San Diego** researchers at the **San Diego Supercomputer Center (SDSC)** deployed an expert networking and data, models, and services workflow platform through a community-driven initiative called the Wildfire Commons. The national-scale effort advances science, technology, and AI solutions to wildland fire challenges in an era of more frequent and devastating megafires.

- By leveraging simulations powered by **HPE** supercomputers at the **Argonne Leadership Computing Facility (ALCF)**, **TAE Technologies** is tackling the critical challenges of stabilizing high-density, high-pressure plasma and designing reactors capable of sustaining fusion, bringing the vision of near-limitless clean energy closer to reality.

- **The U.S. Department of Energy's National Nuclear Security Administration** collaborated with **Cornelis Networks** for the Next-Generation High Performance Computing Network (NG-HPCN) project. Together, they deployed Cornelis Networks' CN5000 to power the 952-node "Lynx" cluster at **Lawrence Livermore National Laboratory (LLNL)** with the objective of driving national nuclear and energy security as part of the Commodity Technology Systems (CTS-2) network infrastructure refresh, in partnership with **Dell Technologies**.

- An alliance between **Novo Nordisk**, **Columbia University**, **OpenFold**, and **Amazon Web Services (AWS)** offers faster, cheaper protein modeling speeds and democratizes state-of-the-art methods for the bio community. The consortium halves alignment runtime/cost and allows researchers to scale training to 256 GPUs.

- **The National Energy Research Scientific Computing Center (NERSC)** and **VAST** have come together with the Doudna project, which transforms storage from a static repository into a dynamic, workload-intelligent fabric that seamlessly adapts to the real-time, data-intensive demands of next-generation science. The effort is expected to accelerate breakthroughs across fields like fusion energy, genomics, astrophysics, and quantum simulation by enabling real-time, data-driven science at unprecedented scale.

17. Top HPC-Enabled Scientific Achievement

- **Duke University** researchers developed the first framework to capture 3D coronary blood flow over millions of heartbeats by combining wearable data, clustering algorithms, and heterogeneous HPC/cloud platforms. This transformed cardiovascular digital twins from single-heartbeat snapshots into longitudinal disease-monitoring tools.

- By uniting exascale computing, algorithmic innovation, and world-first scientific capability, ORBIT-2 redefines the frontier of climate science and demonstrates how HPC can directly advance humanity's ability to understand and prepare for a changing planet. To that end, researchers at **Oak Ridge National Laboratory (ORNL)** trained ORBIT-2 on the **Frontier** supercomputer, leveraging 65,536 GPUs and sustaining 4.1 exaFLOPS of throughput with 92% to 98% parallel efficiency.

- A research team focused on better understanding how magnetic fields influence the highly turbulent motions through the interstellar medium (ISM). Using HPC resources at the **Leibniz Supercomputing Centre (LRZ)**, the team was able to model turbulence in the ISM in unprecedented detail, calling long-held assumptions of the role of magnetic turbulence into question in the process and providing new research directions for next-generation experiments in space.

- **UC San Diego**'s WIFIRE program helped responders fight the January 2025 fires in Los Angeles and aims to mitigate the risk of future megafires. The WIFIRE program informs first responders with real-time models of how a wildfire will spread, creating predictive maps out of data points such as wind conditions, topography, temperature, and the dryness, density, and types of foliage that serve as fuel.

- BEAST, running on the **Pittsburgh Supercomputing Center (PSC)**'s Bridges-2, analyzed Hubble Space Telescope light spectra from 1.5 million distinct stars in the Magellanic Cloud galaxies all at once. This work agrees well with earlier, smaller-scale studies, also offering the possibility of recognizing in fine detail how these galaxies evolved and how they're both similar to and different from the Milky Way.

- Researchers from the **University of Illinois** used **National Center for Supercomputing Applications (NCSA)** resources to perform a detailed analysis of DNN-based top quark taggers using neural activation patterns and other new Explainable AI tools they had developed. This led to the creation of the Particle Flow Interaction Network, a powerful new AI model that has outperformed previous tools. This work on Evidential Deep Learning led to more performant, interpretable, and model-independent anomaly detection for future scientific discovery.

- The climate model ICON is the central research tool at the **Max Planck Institute for Meteorology (MPI-M)**. The ICON team presented the first-ever global simulation of the full Earth system at 1.25 km grid spacing, achieving the highest time compression with an unseen number of degrees of freedom. To achieve this landmark simulation, they harnessed the full power of the GH200 Superchip on Alps and Jupiter, and used **NVIDIA** Grace CPUs and Hopper GPUs to carefully balance Earth's components in a heterogeneous setup, utilizing optimizing acceleration techniques available in ICON's codebase.

18. Top Energy-Efficient HPC Achievements

- **VDURA**'s NVMe flash + capacity-tier architecture delivers the industry's best power efficiency, whether you measure in $/kW or TB/kW. It puts flash only where speed matters and relies on energy-lean HDDs

for the rest. That smart balance slashes cost per terabyte by about 50 % and drives 44 % more usable capacity per watt.

- **Xcimer Energy**'s IrisNet is a deep learning model trained on extensive HPC simulations run on power-efficient **Amazon Web Services (AWS)** Graviton4 processors, enabling rapid optimization of laser beamlines for inertial confinement fusion while achieving significant energy savings over traditional x86-based computing and full-scale simulations. This integration of advanced simulations and AI stands out for its potential to accelerate the commercialization of clean, carbon-free energy through enhanced power efficiency.
- **The Simons Foundation** is at the forefront of cutting-edge research. It uses direct liquid cooling for multiple clusters in its research, which allows it to provide sustainable compute that ensures the foundation's mission and focus are interwoven into practical applications, while delivering maximum energy efficiency.
- **NVIDIA**'s new software technology in cuBLAS doubles the energy efficiency score of any system using the NVIDIA Blackwell B200 accelerator. Specifically, cuBLAS is a GPU-accelerated library that implements the Basic Linear Algebra Subprograms (BLAS) standard for HPC and AI applications on NVIDIA's GPUs.
- **Amazon Web Services (AWS)** demonstrated how Graviton-accelerated Multiple Sequence Alignments (MSAs) enable researchers utilizing advanced protein structure prediction model OpenFold3 to perform cost-aware training. Switching MSA generation to Graviton4 r8g cut runtime by about 50% and cost by about 55%. Training realized an approximately 85% cost reduction using Capacity Blocks + Spot, materially lowering energy per prediction.
- **HPE** AI Mod POD is the first native direct liquid-cooled modular data center able to run extreme density HPC and AI systems. It supports up to 400kW per rack, with integrated heat recovery to help organizations achieve their rapidly expanding AI needs, power challenges, and sustainability targets.

19. Top Supercomputing Achievement

- **Lawrence Livermore National Laboratory (LLNL)** launched El Capitan, the world's fastest supercomputer (1.742 exaFLOPS HPL, 2.79 exaFLOPS peak) and the first exascale system dedicated to national security. LLNL's El Capitan achieved a rare "triple crown", ranking #1 on TOP500 (HPL), HPCG, and HPL-MxP (mixed-precision/AI), demonstrating leadership across real-world and AI workloads. Built with HPE's Cray EX and AMD Instinct MI300A APUs, it delivers massive scale (11M+ cores, ~44K APUs) with strong energy efficiency (58.9 GF/W).
- Elon Musk's Colossus AI infrastructure, said to be one of the most powerful AI computing clusters in the world, was designed to push the boundaries of AI. The massive computing system consists of 200,000 GPUs, all running on **Tesla** Megapack batteries.
- **Opera** is at the forefront of developing user-centric web browsers that deliver the best online experience. To support the most advanced AI models and facilitate generative AI applications, Opera required a high-density supercomputing cluster, and conducted a global search for a site with available high-density capacity, energy-efficient, attractive power pricing, and low environmental impact, in a

short timeframe. By partnering with **atNorth**, Opera's Iceland KEF-1 supercomputer was operational in just six months.

- The **Friedrich-Alexander-Universität (FAU)**'s Helma supercomputer achieved #3 worldwide in the IO500 Production list and #1 among Lustre-based systems. It uses an open-source Lustre + xiRAID storage system built entirely from commodity NVMe hardware to provide fast data to the largest University AI cluster within Germany.

- JUPITER is Europe's first exascale supercomputer. Designed by the **Jülich Supercomputing Centre (JSC) at Forschungszentrum Jülich** in collaboration with **EuroHPC JU** and procured by EuroHPC Joint Undertaking, JUPITER stands as the first European supercomputer capable of performing one exaflop. It is also the most powerful system in Europe, combining outstanding performance with exceptional energy efficiency.

- Researchers at **Oak Ridge National Laboratory (ORNL)** trained ORBIT-2 at unprecedented scale on the Frontier supercomputer, utilizing 65,536 GPUs and sustaining 4.1 exaFLOPS of performance with 92–98% parallel efficiency. This advance has enabled training at resolutions never before possible. ORBIT-2 also introduces novel architectural innovations (Reslim ViT and TILES) that break through long-standing scalability limits in AI models. By combining supercomputing scale, algorithmic innovation, and real-world impact, ORBIT-2 represents a landmark milestone at the frontier of AI and climate science.

- Anticipating the explosion of data, analytics, and the combination of modeling and simulation with AI, **HPE** delivered the first, second, and third verified exascale-class supercomputers, solving the intricate needs of power and cooling, hardware, software, storage, and interconnect, in its quest to fulfill the DoE's mission of next-generation science and discovery. The systems were installed at **Oak Ridge National Laboratory (ORNL)**, **Argonne National Laboratory (ANL)**, and **Lawrence Livermore National Laboratory (LLNL)**.

* 20. Top 5 New Products or Technologies to Watch *(Please pick 5)*

- [ ] ALCHEMI
- [ ] CUDA-Q
- [ ] NVIDIA Quantum-X800 InfiniBand Platform
- [ ] Phison Pascari X-Series Enterprise SSDs
- [ ] Ultra Ethernet Consortium-complaint Etherent
- [ ] Parallel Works ACTIVATE
- [ ] Penguin Solutions OriginAI® AI Infrastructure Solution
- [ ] NVIDIA Spectrum-X Ethernet Platform
- [ ] Mediaflux® for AI
- [ ] Hammerspace Data Platform

- [ ] AI Modular Data Center AI POD
- [ ] GigaIO SuperNODE™ — Composable Scale-Up GPU Server
- [ ] xiRAID Classic
- [ ] EC2 P6e-GB200 UltraServers
- [ ] NextSilicon Maverick-2
- [ ] Lenovo ThinkSystem SC777 V4
- [ ] GigaIO FabreX — Composable PCIe Gen5 Interconnect Fabric
- [ ] AWS Parallel Computing Service Managed Accounting
- [ ] NVIDIA Photonics
- [ ] YellowDog Cloud Compute Scaling Platform
- [ ] Amazon Nova – multimodal FMs on AWS
- [ ] VDURA Data Platform
- [ ] Distributed Asynchronous Object Storage (DAOS)
- [ ] Cornelis CN5000
- [ ] NVIDIA Warp
- [ ] NVLink Fusion
- [ ] VAST Data AgentEngine
- [ ] Warewulf Pro

* 21. Top 5 Vendors to Watch *(Please pick 5)*

- [ ] Altair
- [ ] Amazon Web Services (AWS)
- [ ] AMD
- [ ] Architecta
- [ ] Broadcom
- [ ] Cerebras Systems
- [ ] DDN
- [ ] Dell Technologies

- [ ] Google
- [ ] HPE
- [ ] IBM
- [ ] Intel
- [ ] Lenovo
- [ ] Microsoft
- [ ] NetApp
- [ ] NextSilicon
- [ ] NVIDIA
- [ ] Penguin Solutions
- [ ] SambaNova
- [ ] Starfish Storage
- [ ] Supermicro
- [ ] Swiss Vault
- [ ] ThinkParQ (the company behind BeeGFS)
- [ ] VAST Data
- [ ] VDURA
- [ ] Xinnor

22. Workforce Diversity & Inclusion Leadership Award

- ◯ Over the last three years, the **Research Computing Center at Florida State University (FSU)** has created and run the interdisciplinary data humanities initiative to foster and educate about the use of HPC in the humanities, social sciences, and the arts. During this time, it installed two new tools (Open Refine and Agisoft Metashape) into the open on-demand infrastructure, supported six grant applications with two subsequent long-term projects, and facilitated over 30 workshops at our institution, impacting over 325 students and faculty.

- ◯ **STEM-Trek** is a nonprofit whose mission is to advance diversity and inclusion in the HPC workforce by providing travel support, mentoring, and advanced skills training to STEM scholars from underrepresented and underserved backgrounds. Over the past year, STEM-Trek has enabled diverse scholars from across the globe to attend major HPC conferences and training events, paired early-career researchers with experienced mentors, and fostered connections between HPC practitioners and veterans, displaced workers, and other underrepresented groups.

- **VDURA** is committed to creating lasting change in diversity and inclusion, both within its organization and across the HPC community. Externally, the company has sponsored multiple **Women in High Performance Computing (WHPC)** events and, in collaboration with another supercomputing center, a VDURA employee co-founded the Pittsburgh WHPC chapter. Internally, VDURA cultivates an environment where women can succeed through programs like VDURA Coffee Chats, which provide a safe space to share experiences, seek guidance, and gain executive support.

- **Amazon Web Services (AWS)** is a corporate champion of **Women in HPC (WHPC)**, driving global mentorship, training, and visibility programs (including SC/ISC workshops and fellowships) that measurably expand women's participation and leadership pipelines in HPC.

- The **National Center for Supercomputing Applications (NCSA)**'s The Future of Discovery: Training Students to Build and Apply Open Source Machine Learning Models and Tools (FoDOMMaT) Research Experiences for Undergraduates (REU) program fosters workforce diversity and inclusion by training and mentoring undergraduate students from diverse academic backgrounds in AI and HPC. This enables them to contribute impactful, openly shared research at the forefront of science and engineering. The program is perfect for students who wish to explore research but lack access to such opportunities, including those from two- and four-year colleges where research resources and specialized training may be limited.

- **bp** allows women to thrive in their roles, and in turn provides much value to their business. A few great examples are Elizabeth L'heureux, who is their HPC Director, Arianna Martin, who is one of their HPC Engineering Administrators, and Muhong Zhou, who is a Senior HPC Software Engineer. These women play a big role in keeping the Centre for High Performance Computing running smoothly and successfully, which is a big contributor to bp's success. In addition, bp has been consistently recognized for its leadership in diversity and inclusion, earning multiple awards and recognition over the last decade as a leader in LGBTQ employee rights.

23. Outstanding Leadership in HPC (Please select top choice and add a 2nd choice in comment field below)

- **Katie Antypas, National Science Foundation**

  Katie Antypas has led the National Artificial Intelligence Research Resource (NAIRR) Pilot at the National Science Foundation (NSF), effectively bringing together industry, government, and academia to support AI researchers across the country and demonstrating how a national AI research resource could work to accelerate AI research.

- **Demis Hassabis, Google DeepMind**

  Demis Hassabis is the co-founder and CEO of Google DeepMind, one of the world's leading AI research groups. DeepMind has produced landmark research breakthroughs such as AlphaGo, the first program to beat the world champion at the game of Go, and AlphaFold, which was heralded as a solution to the 50-year grand challenge of protein folding.

- **Rick Stevens, Argonne**

Rick Stevens is Argonne National Laboratory (ANL)'s Associate Laboratory Director for the Computing, Environment and Life Sciences (CELS) Directorate and an Argonne Distinguished Fellow. He is responsible for ongoing research in the computational and computer sciences, from HPC architecture to the development of tools and methods for bioinformatics, cancer, and infectious disease, among others, in science and engineering. Stevens has significant responsibility in delivering on the national initiative for exascale computing and developing the national initiative in AI.

◯ **Charlie Catlett, Trillion Parameter Consortium**

Charlie Catlett is a Senior Computer Scientist at Argonne National Laboratory (ANL), a visiting Senior Scientist at the Department of Computer Science, The University of Chicago, and the Executive Director of the Trillion Parameter Consortium (TPC.dev). The overarching focus of his work at the Consortium is to bring together groups interested in building, training, and using large-scale models with those who are building and operating large-scale computing systems.

◯ **Elizabeth L'heureux, bp**

Elizabeth L'heureux, a successful engineer and biophysicist at bp, leads a strong HPC team known for embracing cutting-edge technology and managing complex systems. To date, bp is currently quadrupling the processing power of the HPC Centre at their US headquarters in Houston, as they upgrade seismic imaging and analysis for exploration.

◯ **Haley Sharon Carter, Northwestern University**

Haley Sharon Carter has become Northwestern University's most valuable computational resource, helping Northwestern researchers succeed with complex computational projects by applying her technical expertise in HPC optimization, along with her patient, approachable support that makes advanced bioinformatics accessible to everyone.

◯ **Satoshi Matsuoka, RIKEN**

Satoshi Matsuoka has been the director of RIKEN Center for Computational Science (R-CCS) since 2018. He is responsible for developing the supercomputer Fugaku, which has become the fastest supercomputer in the world in all four major supercomputer rankings in 2020 and 2021 (TOP500, HPCG, HPL-AI, Graph500). Matsuoka oversees a multitude of ongoing cutting-edge HPC research, including investigating Post-Moore era computing and the future FugakuNEXT supercomputer.

◯ **Trish Damkroger, HPE**

Trish Damkroger is SVP and Chief Product Officer of HPC & AI at HPE. She leads more than 2,000 engineers, supply-chain experts, and service specialists who design and deliver some of the world's largest computing systems. Since joining HPE in 2022 after senior roles at Intel and Lawrence Livermore National Laboratory, (LLNL), Damkroger has steered the unit's end-to-end product strategy, extending exascale technologies into enterprise-ready platforms and supporting customers that range from the world's largest research labs to startups.

◯ **Devesh Tiwari, Northeastern University**

Professor Devesh Tiwari is an internationally recognized leader in HPC and quantum computing whose transformative efforts this year, ranging from deploying Northeastern University's cutting-edge Explorer Cluster with NVIDIA H200 GPUs to implementing community-informed mentorship and sustainable computing practices, have redefined how academic HPC personnel are trained and HPC resources are shared and scaled.

◯ **Amanda Hassenplug, Purdue University**

Amanda Hassenplug has been cited for her exceptional management of the ANVIL Research Experience for Undergraduates (REU) program, where her mentorship and professional development efforts led to every student receiving full-time job offers at national laboratories, and for her service on the PEARC25 Student Program Committee, advancing HPC education and student engagement nationwide.

◯ **Thierry Pellegrino, AWS**

Thierry Pellegrino is the Global Head of Advanced Computing at Amazon Web Services (AWS), overseeing HPC, domain-specific ML, IOT, and Quantum for the company. He is mobilizing cloud + on-prem hybrid advanced computing at a global scale, catalyzing industry/academic collaborations that push HPC and AI frontiers.

◯ **Dr. Bill Nitzberg**

Dr. Bill Nitzberg is the CTO, PBS Works at Altair. He has dedicated his life to advancing HPC and is celebrated not only for his technical leadership but also as a true character whose presence has shaped and inspired the community for decades.

* 24. Thanks for your participation!

PLEASE NOTE: We require a name, organization, and email address in order to help prevent ballot box stuffing. Your responses will remain anonymous.

Full Name

Organization

Email Address

Done >>