



## ALTAIR HPC SUMMIT 2021 AGENDA

Tuesday, May 11

### Keynote: Altair, CEO & Founder & HPCwire “Person to Watch” for 2021 James R. Scapa

- **Speaker:** James R. Scapa, Chairman, Founder & CEO
- **Abstract:** In this keynote, Altair CEO and Founder James Scapa addresses how HPC has evolved over the past 30+ years and how Altair has been putting the pieces in place to drive the future of multi-dimensional HPC. The presentation outlines Altair’s vision for the convergence of HPC, artificial intelligence, and simulation.

### Keynote: Intel, Redefine What's Possible with HPC

- **Speaker:** Trish Damkroger, Vice President and General Manager, HPC, Data Platforms Group
- **Abstract:** We are seeing a mass increase in data, coming at us from all directions. Whether you are running AI-enhanced HPC applications on-prem or HPC workloads in the cloud, you need platforms and solutions that will offer the performance, scalability, and flexibility to redefine what’s possible with HPC.

### Feature Presentation: Aston Martin Cognizant Formula One Team, On and Off the Track, Every Second Counts: the Critical Role of HPC Optimization

- **Speaker:** Laurence Sawbridge-Praties, HPC Engineer
- **Abstract:** This presentation will discuss the role of Altair® Grid Engine®, primarily within the CFD department and other areas of the team which have seen its benefits and have integrated it into their workflows. All the departments using Grid Engine work directly on the development of the car. Its proven contribution to overall performance and Aston Martin’s position on the Formula 1 grid will be demonstrated.

### Keynote: Oracle, Artificial Intelligence Taking to the Cloud

**Speaker:** Clay Magouyrk, Executive Vice President for Oracle Cloud Infrastructure (OCI)

**Abstract:** In this presentation, Clay Magouyrk outlines Oracle's vision the future of artificial intelligence, including how AI and automation will drive the next wave of business transformation, the proliferation of technologies focused on understanding human language, intent, and context, like natural language processing (NLP), and how they are reshaping our interactions and businesses, as well as the critical role of high-performing, low-latency infrastructure in driving successful AI initiatives.

## **Feature Presentation: National Center for Atmospheric Research, Job Scheduling on NCAR's Next-generation Supercomputer**

**Speakers:** Irfan Elahi, Brian Vanderwende, John Blaas

**Abstract:** NCAR provides the atmospheric research community with the shared resources necessary to work on the most important scientific problems of the day. Today, NCAR offers scientists a host of tools, from community models, to curated observational data services, to research aircraft and observational facilities, to computational and data resources, to education and outreach. NCAR's in-house scientists collaborate with their peers in academia, government, and the private sector to build, refine, and expand NCAR's community resources in order to make them as relevant and useful as possible.

Through its High-Performance Computing Division (HPCD), NCAR's Computational and Information Systems Laboratory (CISL) supports and advances this research in the geosciences with the world-class high-performance computing and high-performance storage needed for the execution of large, long-running numerical simulations, data archiving, a full-featured computing software tools environment, expertise, and support services that emphasize user productivity and cost-effectiveness. CISL's high-performance computing and storage resources are architected to provide performance and capacity commensurate with its computational resources, along with high-speed networking and data communication capabilities tailored to the requirements of a national and international scientific community.

For workload management and resource scheduling on its current flagship supercomputer, Cheyenne, and on its data analysis and visualization/machine learning (DAV/ML) cluster, Casper, CISL uses Altair® PBS Professional®. Both of these current systems and NCAR's new NWSC-3 supercomputer, Derecho, will exploit features like cloud bursting, fairshare, power-user and maintenance reservations, resource assignment with control groups, high-throughput hierarchical scheduling, green provisioning, and energy-aware scheduling. The presenters will describe how CISL utilizes the features of PBS Professional to maximize resource utilization while balancing resource usage across the diverse user community within a hierarchical allocation matrix, from large to small projects and thousands of users. They will also address CISL's plans for future use of Altair's Accelerator™ Plus, GUI interface, and additional HPC features not currently implemented on Cheyenne.

## **Day One Breakout Sessions**

### **Breakout: Life Sciences feat. Johnson & Johnson, Bristol Myers Squibb, Seqera Labs**

**Panelists:** Brian Wong, DCM, Research Computing at Bristol Myers Squibb, Satish Murthy, Senior Manager, Janssen Pharmaceutical Companies of Johnson & Johnson, Robert Lalonde, Chief Commercial Officer, Seqera Labs

**Hosts:** William Bryce, VP Product, Altair; Dr. Rosemary Francis, Chief Scientist, Altair

**About the Session:** How do life sciences organizations ensure scientists and researchers have access to the critical problem-solving power they need? In this breakout session, HPC experts from Johnson & Johnson, Bristol Myers Squibb, and Nextflow developers Seqera Labs will address the unique time, budget, and uptime requirements faced by life sciences organizations and the tools they use to optimize the HPC environments that make real-world discovery possible.

Agenda: Introduction to the panel and HPC for life sciences, guided discussion: workload requirements and challenges, technology trends, real-world examples, live Q&A: open to all participants, ask questions live and engage in discussion with panelists and peers.

### Who Should Attend?

Life sciences industry professionals, HPC professionals, and end users interested in workload management strategies for high-volume datasets, effective data lifecycle management, and workload optimization including cloud migration and automation.

### Breakout: Compute Challenges in Semiconductor Design

**Hosts:** Stuart Taylor, VP Product Development, Altair, Dr. Rosemary Francis, Chief Scientist, Altair, Brian Janes, Senior Director, High-performance Computing, Altair

**About the Session:** In this session, panelists will discuss how organizations leading the highly competitive semiconductor design space ensure their compute infrastructure is tuned to optimize the high-intensity workloads that are producing faster, smaller, and more innovative chip designs. Panelists will cover recent changes in tools like Altair Accelerator including its streaming data service, motivation and implementation choices, Kafka, and building a data pipeline. They will also discuss direct drive, cloud deployment including Rapid Scaling and storage- and I/O-aware scheduling, and hybrid cloud topics such as license sharing and I/O profiling to assess and curate jobs for cloud use.

**Who Should Attend:** HPC and design engineering professionals with a special interest in demanding, high-throughput workloads such as those of the semiconductor industry.

## Wednesday, May 12

### Featured Roundtable: AI Takes to the Clouds, Google Cloud, Oracle, Microsoft Azure, Intel, NVIDIA

**Panelists:** Dr. Vikram Saletore, Deep Learning Performance, Intel, Dr. Bill Magro, Chief Technologist, High Performance Computing, Google, Nidhi Chappell, GM, Azure HPC, AI, SAP and Confidential Computing Initiative, Karan Batta, Vice President of Product, Oracle, Shanker Trivedi, Senior Vice President, Enterprise Business at NVIDIA

**Abstract:** Gartner predicts that AI will become more operationalized by 2024. And, as organizations increasingly embark on their digital transformation journey, AI, data-driven analytics, and the convergence in the cloud are the centers of this transformation. AI is just the technology to revolutionize cloud computing in making intuitive, connected experiences possible.

As companies find more ways to deploy AI for real-time predicting and prescribing, the resource requirement for training data-heavy models will yield increased demand for HPC infrastructure in the cloud. Additionally, organizations are exploring the application of AI to HPC, using AI to augment HPC optimization and even automate cloud migration.

In this panel we speak with chip industry leaders and cloud service providers to get their take on the impact and trends they are seeing from companies who are taking an AI-first approach and how it's driving the move to the cloud.

### Featured Industry Update: Hyperion Research, HPC Market Update and Survey Findings

**Speaker:** Alex Norton, Principal Technology Analyst and Data Analysis Manager

**Abstract:** HPC data centers have shifted their priorities to not only time-to-solution for applications, but also time-to-solution for a given cost. In this session, Hyperion Research will reveal survey results of a study targeted at HPC end users to help support the shifting trend of focus for HPC datacenter designs.

The study focused on key aspects of HPC datacenters, including processor and accelerator trends and adoption rates, HPC middleware and software, HPC and AI application trends, HPC cloud adoption, and critical components of future datacenter deployments. Hyperion Research will share findings such as how AI adoption in HPC continues to grow at a high rate, both from an investment perspective in dedicated server infrastructure and an adoption perspective, and how systems are also incorporating new memory and interconnect technologies to increase application performance and time-to-solution.

### **Featured Presentation: Safran Seats**

**Speaker:** Daniel Boast, VP Stress Methods

**Abstract:** Daniel Boast of Safran Seats will outline how his team accelerates the product design and development process with a turnkey, virtual Altair HPC appliance: Altair Unlimited™.

### **Featured Presentation: Microsoft Azure, Using I/O Profiling to Migrate and Right-size EDA Workloads in Microsoft Azure**

**Speaker:** Michael Requa, Sr. Program Manager, Microsoft

**Abstract:** Semiconductor companies are taking advantage of Microsoft Azure HPC infrastructures for their complex electronic design automation (EDA) software. When one of the largest semiconductor companies asked for help using Azure to run its EDA workloads, Microsoft teamed up with Altair.

The semiconductor company was running a relatively large design of 100 million transistors. We used Altair Breeze™, an I/O profiling tool, to troubleshoot and tune the company's workload before and after migrating it to Azure. Breeze revealed I/O patterns that had previously gone unnoticed and showed us how to significantly improve application run time.

This presentation will outline how Microsoft used Breeze to diagnose I/O patterns, choose the workflow segments best suited for the cloud, and right-size the Azure infrastructure. The result was better performance and lower costs for our semiconductor customer.

### **Feature Presentation: National Institute for Aviation Research (NIAR) and Oracle, Scalability Study of a Large Ditching Model Using Radioss Mono and Multi-Domain Techniques**

**Speakers:** Rafael Bini Leite, Research Associate Computational Mechanics Lab

**Abstract:** The scalability performance of Altair Radioss™ software when running in parallel with up to 1152 cores, using mono and multi-domain techniques, was assessed for a large ditching model. The ditching model comprises a simplified cylindrical fuselage impacting on a flat surface fluid domain (water). The water is modeled using a hybrid approach, wherein the region of the impact the fluid domain is modeled with smoothed particle hydrodynamics (SPH) and the surrounding regions are modeled with Lagrangian solid elements. Within the range of cores tested for the multi-domain (72 to 288 cores), the overall efficiency of the multi-domain simulations is significantly higher than the monodomain, taking almost a third of the time to complete the simulation when compared to the monodomain model for 72 cores. Both techniques have shown excellent scalability, with high scaling efficiency and with outstanding reliability of the results.

## Partner Breakout Sessions

### Partner Breakout: Intel, Innovations for the 'Data Era'

**Speaker:** Kirti Devi, Director of HPC Global Alliances, Data Platforms Group, Intel

**About the Session:** This is the era of data. Advances in today's technological front with HPC infrastructure, artificial intelligence, cloudification, and data analytics are coalescing to provide a source of key competitive differentiation to organizations uniquely positioned to leverage this data influx strategically, converting it from merely information to insights. This presentation will highlight the innovations in the industry and at Intel with its portfolio of products, technologies that will redefine the world of compute.

### Partner Breakout: Microsoft, Altair nanoFluidX on Azure's new NDv4

**About the Session:** In this session Microsoft will showcase Azure's new ND A100 v4 VM series running Altair nanoFluidX™. By leveraging Azure's most powerful and massively scalable AI VM, available on demand from eight to thousands of interconnected Nvidia GPUs across hundreds of VMs, users can scale nanoFluidX to over 100M particles with ease.

**Speaker:** Jon Shelley, Principal PM Manager, Microsoft Azure

### Partner Breakout: Oracle, HPC with Oracle Cloud: Bare Metal Computing in a Cloud Environment

**About the Session:** TBC

**Speaker:** Kevin Jorissen, Distinguished HPC Cloud Architect with Oracle Cloud