



Speaking Engagements in the SC21 Hubb

Monday, November 15th

| <i>Session Title</i> | <i>Description</i> | <i>Time (CST)</i> | <i>Session Type</i> |
|---|--|-------------------|-----------------------------------|
| AI and Data Science on Workstations: From Hype to Reality Across Industries | Dell Precision Data Science Workstations deliver industry-leading solutions, not just for developing AI/ML/DL models, but also deploying those models to solve hard industry problems. We will demonstrate several examples of industry solutions developed and deployed on Dell Precision Data Science Workstations. | 10:00 AM | Breakout Session |
| Supercomputing Wales | Supercomputing Wales is the national supercomputing infrastructure for Wales, UK. A consortium of four universities has come together to offer high-performance computing and research software engineering to researchers working in a range of scientific fields including earth observation, catalysis, epidemiology and gravity exploration. Join us to learn about why we do what we do, our HPC facilities, and how we've been helping the fight against the pandemic. | 10:00 AM | Exhibitor Session |
| Density-Optimized Multi-node Server for Extreme Compute at Scale | The Dell EMC PowerEdge C6520 is the ideal server designed to support dynamic compute-at-scale requirements in a dense form factor. The PowerEdge C6520 maximizes performance per U by supporting 4 dual-socket compute nodes in a 2U chassis with air and direct-liquid-cooling options to power the highest performing processors. | 10:30 AM | Breakout Session |
| Design, Build, Test and Help Your Vision Come to Life | Experience our end-to-end solutions when you set up an engagement with us, designed to solve your business challenges and optimize innovation within your organization. Work with Dell Technologies experts in our dedicated labs or remotely from any location with the latest showcases of our offerings. | 11:00 AM | Breakout Session |
| Bring the Edge to your Organization with Dell EMC Modular Data Center Solutions | Dell EMC micro Modular Data Centers are self-contained units that can feature power, cooling, gateways and can host up to several racks of IT. They are designed with the security, environmental and performance capabilities to be placed in remote locations where real estate space is limited, like at the base of cell towers, parking lots, or in remote locations. | 11:30 AM | Breakout Session |
| Geekfest! | Visit with members of the Dell Technologies Chief Technology Office about the trends they're seeing and the new technologies they're exploring to help you get ready for what's next. | 12:00 PM | Breakout Session |
| Enabling Present and Future AI Models | Graphcloud is a MK2 IPU-POD scale-out cluster, offering a simple way to add state of the art machine intelligence compute on demand, without the need for on-premise hardware deployment. Graphcloud is ideal as you scale from experimentation and proof of concept projects to pilots and production systems. | 12:30 PM | Breakout Session |
| Going with the Flow - Neural Nets for Fluid Dynamics | Find out how artificial intelligence can stand in for simulation and modeling workloads that typically run in high performance computing systems. In this project, Varun Shankar, Carnegie Mellon University and Dr. Luke Wilson, Dell Technologies share how the use of trained neural networks can allow design teams to quickly evaluate the viability of many different design alternatives in comparison to running computationally-intensive simulations for each alternative under consideration. | 1:00 PM | Exhibitor Session |
| Work Anytime, Anywhere with Virtual Desktop Infrastructure | Large scale systems are increasingly being used to solve complex machine learning (ML) problems in natural language processing, computer vision, drug discovery and recommendation systems. Models with trillions of parameters that need exaflop scale compute and petabytes of memory are not too far out in the future. Attend this session to learn more about how Graphcore is addressing these challenges with Intelligent Processing Unit (IPU), which is a purpose-built accelerator for the most demanding compute and memory bandwidth needs of modern ML models. Our disaggregated approach allows customers to start small and then seamlessly scale to thousands of IPUs to tackle multi-trillion parameter models. | 3:00 PM | Breakout Session |

Tuesday, November 16th

| <i>Session Title</i> | <i>Description</i> | <i>Time (CST)</i> | <i>Session Type</i> |
|---|--|-------------------|-----------------------------------|
| Find Out About Modern, Flexible Scale-Out File Storage | Dell EMC PowerScale, the world's most flexible scale-out NAS solution, is designed to be flexible and reliable at any scale. Regardless of the kind of data, where it lives, or how big it gets, your data lake always stays simple to manage, simple to grow, simple to protect, and simple enough to handle the most demanding workloads of today and tomorrow. | 10:00 AM | Breakout Session |
| Advanced Visualization and Data Analysis of HPC Cluster and User Application Behavior | This work presents cutting-edge visualization, monitoring, and management solutions for HPC systems to understand the status of high-performance computing platforms and provide insight into the interactions among platform components. Benefiting from the greatly increased level of detail available from modern baseboard management controllers through Redfish Telemetry and real-time correlations via API and CLI interfaces to HPC job schedulers, this work provides much greater detail than previous similar projects. | 10:00 AM | Exhibitor Session |
| Hot Topics in the HPC & AI Innovation Lab | The Dell Technologies HPC & AI Innovation Lab has been implementing new hardware, software and cooling systems while engineering solutions. Stop by to learn more about the Innovation Lab, to ask questions and to hear about lessons learned in upgrades. | 10:30 AM | Breakout Session |
| Run AI, HPC and data analytics on the same systems, with greater flexibility | With Omnia, you can compose a unified architecture with multi purpose, balanced nodes to support multiple workloads, and quickly re compose resources to meet demands both now and in the future. Omnia is an open source community project started in the Dell Technologies HPC & AI Innovation Lab and you're invited to download and participate at dellhpc/omnia. | 12:00 PM | Breakout Session |
| Flip the Switch to Open Networking with the new PowerSwitch S5448F-ON | The new S5448F-ON 100/400GbE switch provides high-density 100/400 GbE ports and a broad range of functionality to meet the growing demands of today's data center environment. This innovative, next-generation open networking high-density aggregation switch offers optimum flexibility and cost-effectiveness for demanding compute and storage traffic environments. | 12:30 PM | Breakout Session |
| Utilizing Kubernetes and Open OnDemand to Support Virtual Classroom Labs | Alan Chalker and Trey Dockendorf, part of the Open OnDemand (openondemand.org) development team at the Ohio Supercomputer Center provide an overview of supporting virtual classroom labs that utilize Kubernetes in conjunction with Open OnDemand. This is a technical talk that describes in detail all the various work involved with configuring Kubernetes and the Kyverno policy engine to ensure usability and security features standard to Open OnDemand are maintained. | 2:00 PM | Exhibitor Session |
| Enjoy Simplicity in High-Performance and High-Capacity Storage | Consolidate your storage with the PowerVault Series, optimized for entry-level SAN and DAS environments. The ME4 Series is ideal for supporting workloads such as HPC and more. Designed for versatility, the ME4 Series supports a variety of drive types (including SEDs), multi-protocols, software and a choice of two 2U or a high density 5U base systems. Additionally, PowerVault JBOD expansion enclosures are available for PowerEdge Servers that easily add storage capacity. | 3:00 PM | Breakout Session |

Wednesday, November 17th

| <i>Session Title</i> | <i>Description</i> | <i>Time (CST)</i> | <i>Session Type</i> |
|---|---|-------------------|-----------------------------------|
| UCSD Discovers the Sugary Shield Around SARS-COV2 | A team of researchers at UC San Diego is leveraging the massive computational power of systems at the Texas Advanced Computing Center (TACC) to help expose the secrets of SARS CoV-2. SARS-CoV-2 and other viruses use a spiked sugary cloak to attack the cells of a human host. In essence, glycans trick the human immune system into seeing them as harmless. The breakthrough simulations and modeling conducted by the UC San Diego research team have shown the world what the sugary coating looks like and how it tries to hide itself from human immune systems. These efforts reveal that the glycans prime the coronavirus for infection by changing the shape of their spike protein. Scientists hope this research will add to the arsenal of knowledge needed to defeat the COVID-19 disease. | 10:00 AM | Exhibitor Session |
| Get a Performance Boost with the Latest Server Accelerators | Give your applications a boost to accelerate insight and innovation with PowerEdge server graphics processing units (GPUs), field programmable gate arrays (FPGAs) and Intelligence Processing Units (IPUs). | 12:00 PM | Breakout Session |
| Unlock the Value of Your Data with HPC Storage Solutions | Dell Technologies Validated Designs for HPC Storage are delivered with hardware, software and support from Dell Technologies. HPC & AI Innovation Lab engineers develop and tune each solution based on performance characterizations and best practices to simplify installation and speed time to results. Find out about new/updated solutions for BeeGFS, NFS and PixStor. | 12:30 PM | Breakout Session |
| Leveraging AI Techniques in HPC | Find out the many ways artificial intelligence techniques can be leveraged in HPC. Dr. Luke Wilson shares one example of how the use of trained neural networks can allow design teams to quickly evaluate the viability of many different design alternatives in comparison to running computationally-intensive simulations for each alternative under consideration. | 1:00 PM | Breakout Session |

Wednesday, November 17th

| Session Title | Description | Time (CST) | Session Type |
|---|---|------------|----------------------------------|
| PowerEdge XE Servers: Built for Your Most Demanding Workloads | PowerEdge XE8545 is the purpose-built server for complex, emerging workloads that require high performance. These robust servers are designed to optimize the latest industry technologies. With a XE8545 you can develop, train, and deploy cutting edge machine learning models, accelerate complex high performance computing workloads or host accelerated virtualization services. | 3:00 PM | Breakout Session |

Thursday, November 18th

| Session Title | Description | Time (CST) | Session Type |
|---|--|------------|----------------------------------|
| Innovate at Scale with Challenging and Emerging Workloads | The Dell EMC PowerEdge R750xa is a dualsocket/2U rack server that delivers outstanding performance for the most demanding emerging and intensive GPU workloads. It supports 8 channels/CPU, up to 32 DDR4 DIMMs @ 3200 MT/s DIMM speed. The PowerEdge R750xa also supports PCIe Gen 4 and up to 8 SAS/SATA SSD or NVMe drives. With one platform that supports all of the PCIe GPUs in the PowerEdge portfolio, the PowerEdge R750xa the ideal server for AI-ML/DL training and inferencing, and HPC environments. | 10:00 AM | Breakout Session |
| Creating HPC Clusters-as-Code with Omnia | Omnia is a new open-source project that automates the stand-up and tear-down of software-defined HPC clusters for simulation, artificial intelligence and data analytics using various server, storage and network building blocks. Omnia can compose software-defined clustered solutions on bare-metal and virtual machines, whether they be on-prem, in a co-located datacenter or in the cloud. By taking advantage of standard APIs, automation tools and other existing software projects, Omnia can compose software-defined clusters from various infrastructure providers. In this talk we will discuss Omnia's current architecture and capabilities, describe how it composes and instantiates logical clusters in a software-defined manner, and how it can quickly tear down those clusters to repurpose underlying nodes for other cluster deployments. We will then discuss Omnia's future plans, expected features, and how people and organizations can join and contribute to the Omnia community. | 11:00 AM | Exhibitor Forum |
| Overcome Obstacles and Unlock Opportunity with Dell Technologies Services | Make Dell Technologies Services your partner for the digital future. From the edge, to the core, to the cloud, our industry experts offer strategic guidance and proven practical capabilities to help you accelerate time to value of your objectives. | 12:00 PM | Breakout Session |
| Engineering Better Performance in CAE with Intel's Latest Generation Processors | Find out how you can improve the performance of your computer aided engineering (CAE) applications with Dell Technologies! With more than 25 years of experience in the field, Dr. Feyereisen is available to provide technical guidance on all major aspects of solutions including hardware, software and system performance. | 12:30 PM | Breakout Session |
| Advanced IT solutions require smart cooling strategies | Air cooling has long been used effectively for systems with processors using up to 150 watts. With HPC, AI and hyperscale data center growth on the rise, relying on air cooling creates challenges in maintaining equipment performance and efficiency. Today's powerful HPC systems often leverage processors that run at or above 250 watts, bringing liquid cooling to the forefront as a more efficient and cost effective way to cool these high speed processors. | 3:00 PM | Breakout Session |