

IBM Seminars at SC16

Salt Lake Marriott Downtown at City Creek

Tuesday, November 15, 2016

10:30 am - 11:30 am

SEM01- Faster Deployment, Performance with A Reference Architecture for Deep Learning

Abstract: Learn how the IBM Reference Architecture for MLDL (Machine Learning/Deep Learning) helps organizations get faster time to results with popular deep learning frameworks, the investments IBM is making in the open source deep learning ecosystem, and how POWER8 and NVIDIA Tesla P100 GPUs connected by NVLink are opening up new opportunities for insight in this fast growing HPC segment.

Location – Grand Ballroom Salon E

12:00 pm - 1:00 pm

SEM04- Meeting business demands for big data analytics: a practical approach to Apache Spark

Abstract: Apache Spark is emerging as a key application framework for big data processing and complex analytics on both Wall Street and Main Street. It can give users a compelling competitive advantage in analyzing their big data thanks to its speed, its powerful, easy-to-use application programming interfaces (APIs), and a rich set of high-level tools for big data analytics. While Spark has the potential to deliver significant advantages, implementing the open source framework and getting it to perform consistently in a multi-tenant environment requires overcoming several key challenges, including: meeting the performance demands of diverse applications and users; supporting agile and dynamic resource management; increasing management efficiency and consistently meeting SLAs; and enabling expansion to multiple data sources and frameworks. In addition, organizations wonder about being ready for the “next big thing” a couple of years down the road. We will discuss the pros and cons of moving to Apache Spark, whether as part of an existing Hadoop or other environment. We will describe practical solutions for deployment of Spark in a production environment that also provides the flexibility to incorporate additional frameworks and application models as needed, now or in the future. **Lunch will be provided.**

Location - Grand Ballroom Salon F

1:30 pm - 2:30 pm

SEM03- How HPC enables new advances in real time logistics: GPU Accelerated Databases

Abstract: Learn how some of the world’s most challenging analytics problems -- at the intersection of HPC and Data Analytics -- are being solved in real time through the use of accelerated database solutions on IBM Power Systems with NVLink.

Location - Grand Ballroom Salon F

2:45 pm - 3:45 pm

SEM05- Building Better Systems for HPC - POWER8 with NVLink Systems

Abstract: This presentation will introduce the architectural choices behind building better systems for HPC and accelerated computing (especially POWER8 with NVLink), explain what application domains they benefit, and translate this into real-world performance. This is the perfect deep exploration for anyone architecting their next cluster or interested in HPC hardware. We’ll also elaborate how early POWER8 with NVLink installations, combining the best of the OpenPOWER ecosystem including Tesla P100 GPUs, are delivering optimal performance.

Location - Grand Ballroom Salon F

4:00 pm - 5:00 pm

SEM02- Accelerating discoveries in life sciences research and genomics medicine with IBM software-defined infrastructure and architecture

Abstract: Increasing healthcare, drug development, and regulatory costs are driving pharmaceutical, insurance, and healthcare providers toward more collaborative R&D of treatments with a strong focus on genomics medicine. Driven by affordable genome sequencing and IT technologies, researchers and clinicians alike are now seeing extreme data volumes as well as increasing compute and data bottlenecks. Older IT architectures often lack the capacity and agility to support these compute- and data-intensive computing on this scale. In response, IBM developed software-defined infrastructure (SDI) solutions and partnered with industry thought-leaders to create the IBM Reference Architecture for Genomics leveraging SDI. Join IBM and learn how they enable faster data management and access, workload orchestration, and collaboration end-to-end, across sequencing, translational and personalized medicine. This seminar will focus on the three major challenges facing every researcher and clinician: large data ingestion pipelines, dealing with unstructured data, and integrating data from multiple sources for analysis.

Location - Grand Ballroom Salon F

Wednesday, November 16, 2016

8:30 am - 9:30 am

SEM08- From the drawing board to race day, winning with IBM Software Defined solutions

Abstract: Formula One is an incredibly fast moving sport. There can be as little as a few hundredths of a second between first and second place. The speed of innovation is extreme – as many as 100,000 engineering changes are made to each car over the course of a season. And with Formula One regulations becoming more stringent each year, the role played by HPC for virtual analytics and simulation workloads is increasingly important. On the circuit, fast and highly accurate data-driven decision making is critical for success. Join us for breakfast as Nathan Sykes, Head of Numerical Tools and Technologies from Red Bull Racing to discuss the crucial role that IBM Software Defined solutions play to achieve success and how these solutions help – from the drawing board to race day! Learn from IBM subject matter experts about the latest advances in IBM Spectrum Computing and IBM Spectrum Storage and how it is being put to use to win in the real world. **Breakfast will be provided.**

Location - Grand Ballroom Salon F

10:00 am - 11:00 am

SEM07- Challenges and Opportunities: Convergence of HPC, Machine Learning and Big Data

Abstract: A confluence of events is making the convergence of HPC, Machine Learning and Big Data viable for clients using OpenPOWER systems - (i) Enormous data volumes are prohibiting the movement of data between Big Data, Machine Learning and HPC clusters, (ii) Separate organizational cultures maintained by HPC, Machine Learning and Big Data departments are beginning to overlap because of "Data", (iii) The Big Data and HPC communities are coming together to build convergent reference stacks for advancing Science and (iv) The availability of accelerated OpenPOWER systems designed by members of the OpenPOWER community, consisting of members across the HPC, ML and Big Data communities. OpenPOWER clusters can be designed and sized to converge the requirements of HPC, ML and Big Data since OpenPOWER systems have large cache and memory capacity, large memory bandwidth, large IO/network bandwidth and large accelerator IO bandwidth. This allows complex workflows to be developed using a combination of Apache Spark, popular machine learning frameworks and traditional HPC workloads. We will conclude by providing several use cases demonstrating the superior benefits on OpenPOWER systems when converging HPC, Machine Learning and Big Data workflows.

Location - Grand Ballroom Salon F

11:00 am - 12:00 pm

SEM06- HPC Storage – Storage, Tape Systems, File System, Hybrid Cloud and Solutions

Abstract: This session will focus on the challenges and opportunities for addressing data storage and management issues, including discussions on software defined storage, end to end IBM storage solutions for Technical Computing/HPC environments, Spectrum Scale, Elastic Storage Server, HPSS, DCS Storage family, flash storage and tape technologies. In addition, storage tiering and the IBM storage portfolio will be highlighted, which can help reduce data management costs, transforming your organization with better data driven decision in real time using Flash and Storage Technology. Attendees will need to bring signed non-disclosure agreements with them or sign non-disclosure agreements at the door.

Location - Grand Ballroom Salon F

12:15 pm - 1:15 pm

SEM09- Data Centric Cognitive Computing at the Hartree Centre – A Model for Collaboration to Drive Economic Impact

Abstract: On the June 4, 2015 the UK Government announced an IBM and the 'Science & Technology Facilities Council' (STFC) a five year Research & Development Collaboration Agreement (RDCA) aimed at helping to accelerate the capabilities of STFC's Hartree Centre in providing competitive advantage for UK Industry. The IBM Research and the STFC teams have been developing partnerships with UK Companies and other research institutions to this end. This effort is being built on the same OpenPOWER based Data Centric Systems roadmap that will be shipped to the U.S. Department of Energy for the CORAL program. The research collaboration has delivered enhanced cognitive simulations and is continuing to focus on data centric cognitive computing with emphasis on workflows and applications of importance to industry, providing not only data centric systems but a critical mass of computational scientists and engineers to UK industry and research institutions to exploit these systems. This session will provide an overview of collaborative research in progress that is already exploiting data centric systems providing a competitive opportunity for UK industry. **Lunch will be provided.**

Location - Grand Ballroom Salon E

1:30 pm - 2:30 pm

SEM10- New Product Deep Dive: New Uses of Flash with IBM Spectrum Scale and Elastic Storage Server

Abstract: Join us to hear the latest enhancements to IBM Spectrum Scale, Elastic Storage server and Flash. We will introduce new products optimized for HPC and big data analytics; the latest in object storage; and new packaging options. IBM Spectrum Scale is a proven, scalable, high-performance data and file management solution (based upon IBM General Parallel File System or GPFS technology) that's being used extensively across multiple industries worldwide. IBM Elastic Storage Server is a fully integrated server-based storage system for Analytics and HPC, which combines Spectrum Scale software with IBM's latest generation of Power processors, and the latest storage to deliver a complete high-performance, high-value scalable storage solution including erasure coding software, instead of traditional RAID, to deliver high reliability and a graphical user interface to simplify deployment and management.

Location - Grand Ballroom Salon E

2:45 pm - 3:45 pm

SEM11- IBM and Microsoft Azure: Public Cloud Optimized for High-Performance and Analytics Workloads

Abstract: Microsoft and IBM are partnering to support hybrid and stand-alone IBM Spectrum Symphony and IBM Spectrum LSF deployments in Microsoft's Azure public cloud. Join Microsoft and IBM to learn how combining these market leading solutions will provide our customers with the performance and scalability needed for their most demanding workloads. In this session we will review architecture, best practices, and the results of scalability and performance testing.

Location - Grand Ballroom Salon E