

HP-CAST

*HP Consortium for Advanced Scientific and Technical Computing
World-Wide User Group Meeting
HP Servers and HPC Organization*

*July 10th–11th, 2015
HP-CAST 24*

Book of Abstracts *and* *Biographies*

Note: For some presentations no abstract has been received. The abstracts are listed according to the sequence of the agenda version 3.1.

Friday, July 10th, 2015

Hewlett-Packard

Title: HP-CAST 24 Welcome and Overview

Speaker: Frank Baetke, Global HPC Marketing - Academia and Scientific Research

Bio: **Frank Baetke** works for HP's HPC and Hyperscale business within the Industry-standard Servers & Software (ISS) business unit with a special focus on high performance and scale-out systems in academic and research environments. He has a worldwide responsibility and is located in Munich, Germany. Frank Baetke is member of the Board of HP-CAST, the world-wide HPC user group of HP, he is a member of the Advisory Board and a Fellow of the International Supercomputing Conference (ISC) and a member of the Steering Committee of the Indian Supercomputing Conference HiPC. He is also acting as a member of the program committee of several international conferences with a focus on High Performance Computing. He is also representing HP in the European organizations PROSPECT, EESI-2 and EOFS. Prior to joining the world-wide organization, he served as a technology manager in the European Organization of HP. Before its acquisition by Hewlett-Packard in 1995, he held several management positions with Convex Computer Corp. in Richardson, TX. Frank Baetke started his career in 1977 as a researcher and promoted to academic counselor 1979 at the Technical University of Munich. He joined Convex Computer Corp. in 1986 where he built the Central European software support organization and later lead the European marketing organization until the acquisition by HP in 1995. Frank Baetke holds a master degree (Dipl.-Ing.) in engineering and a Ph.D. (Dr.-Ing.) in applied physics from the Technical University of Munich. He has published numerous articles in the field of high-performance computing as well as contributions in related areas. He is a member of GI (German Society of Informatics), the Society of Astronomy, ACM, IEEE and the Max Planck Society.



Abstract:

KIT, Karlsruhe Institute of Technology

Title: Greeting and Welcome from the President of HP-CAST

Speaker: Rudolf Lohner

Hewlett-Packard

Title: HPC-Business Review and Strategic Product and Technology Outlooks

Speaker: Bill Mannell

Bio: Bill Mannell is vice president and general manager of HPC and Big Data Solutions for the HP Servers. The HP Servers portfolio includes HP Apollo, HP Moonshot; HP ProLiant servers, the world's No. 1 server brand; HP BladeSystem, the world's number one bladed infrastructure; and mission-critical computing with the HP Integrity platform. HP Servers power the web and nearly every aspect of daily life from finance and entertainment to healthcare and education.



As the global leader of the HPC and Big Data Solutions team, Bill is responsible for the global end-to-end P&L; and will lead our worldwide strategy and execution in the market's fastest growing category. Bill will also work closely with the Regional Category Managers for Density Optimized (SL/Apollo). Bill joined HP in 2014 and is a seasoned veteran of the servers and high performance computing industry. He joins us from Silicon Graphics International Corp. (SGI), where he was the VP and GM for Compute and Storage Products. He brings with him significant P&L leadership, product marketing and HPC market experience. At SGI, Bill architected the move from proprietary systems into open systems using x86 and Linux and after the merger with Rackable Systems, led product marketing for the combined product families for the service provider market including the biggest brand names in public cloud services. Prior to SGI, Bill worked in the U.S. Government as NASA and the U.S. Air Force, as a lead structural engineer on several leading-edge aircraft and missile programs. Bill holds a Bachelor's degree in Mechanical Engineering and Materials Science from Duke University, Durham, North Carolina. He also holds a Master's of Business Administration from California State University, San Jose, California.

Abstract:

Hewlett-Packard

Title: HPC-Business Review and Strategic Product and Technology Outlooks

Speaker: Scott Misage, General Manager, High Performance Computing HP Server, Hyperscale Business Unit

Bio: Scott Misage is Vice President and General Manager of the High Performance Computing Segment within HP Servers. In this role, Misage manages the worldwide HP business for hardware, software and solutions products for high performance computing, including fabrics, accelerators, storage and HPC software.



After joining the HP high performance computing team in 2001, Misage held several management positions, developing and delivering supercomputing technology and implementations for industrial and academic customer environments. In the last decade Misage and his team have placed five systems among the world's five most powerful computers of their time; including systems in the US, India, Sweden and Japan.

Misage joined HP in 1998 following ten years of experience as an information technology program management specialist with the U.S. Federal Government.

Misage holds a Bachelor of Science degree at Cornell University in Electrical Engineering and a Master of Electrical Engineering also from Cornell University.

Abstract:

Hewlett-Packard

Title: HPC-Business Review and Strategic Product and Technology Outlooks

Speaker: Vineeth Ram

Bio: Mr. Ram has served as Vice President of Global Product Marketing for HP Servers since November 2013.



Previously, Mr. Ram served as the Vice President of Global Marketing for the Avocent Business of Emerson Network Power, a Data Center Solutions division of Emerson (NYSE:EMR) with approx. \$7B in annual revenue, from 2011 to October 2013. Prior to this, Mr. Ram was the Vice President of Global Marketing for Panduit Corp., a Physical Infrastructure Solutions company with approx. \$1B in annual revenue, from 2007 to 2011. Mr. Ram was responsible for all aspects of the global marketing function in both these positions including market research, A/R, PR, Media, strategy & planning, products/solutions, integrated marketing communications, marketing program management, e Before this, Mr. Ram spent approx. 7 years at IBM where he held leadership marketing and strategy roles in the Software, Services and Systems Groups focused on product and solutions marketing & product management. His experience included managing a significant portion of IBM's Enterprise Content Management (ECM) business including acquisitions and, product marketing for pSeries servers. Previously he spent time at Dell, 3M and AMD in a variety of product management, product marketing, strategic marketing and business development roles including P&L management.

Vineeth received a B.E. in Mechanical Engineering from College of Engineering, Guindy, Anna University, Chennai India, a M.S. in Systems Engineering from Wright State University, Dayton OH, a M.B.A. from the University of Texas at Austin, TX and completed post graduate courses at Penn State University, State College, PA.

Mr. Ram's current office location is Houston, Texas

Abstract:

AGH University – ACC Cyfronet

Title: A 1.7 Petaflops Warm-Water-Cooled System: First experience with HP Apollo 8000 at Cyfronet

Speaker: Marek Magrys

Bio: Marek Magrys - HPC Systems Specialist at ACC Cyfronet AGH-UST with over



8 year experience with HPC and Grid systems administration. Mostly focused on storage and parallel file systems and HPC systems architecture. Member of the workgroup responsible for design and procurment of Cyfronet's next-generation supercomputer, Prometheus.

Abstract:

ACC Cyfronet AGH-UST in Cracow, Poland, has deployed a 1.7 PFLOPS HPC system, "Prometheus" in Q1 2015. The talk will describe the procurment and installation process, first experience with the system, it's performance and energy efficiency.

Sauber Motorsport

Title: Using Hybrid HPC-Systems in Formula 1 Racing

Speaker: Philip Postle



Bio: Having spent a lot of time making paper aeroplanes and diverting water on the family farm while growing up, it was not a complete surprise that I ended up studying fluid dynamics at university. While study it became clear that with the rate of advancement in computers, engineering simulation specifically Computational Fluid Dynamic would steadily become a more integral part of the engineering process. Since then my entire career has involved that application of CFD in various industries, including aircraft, motorbikes, road cars, yachts, coffee machines, medical devices and eventually F1. I have been working in F1 now for more than 10 years and 8 of those have been with the Sauber F1 Team in Switzerland.

Abstract: The presentation will give some insight into the background of F1 & the Sauber F1 Team and discuss how HP's understanding of this highly regulated environment lead to the decision to deploy the new Moonshot system at Sauber F1 Team. While also looking into the main applications for which the new computational cluster is being used.

Technical University Dresden and ZIH Dresden

Title: HPC and Big Data: Synergies, Hype and Reality

Speaker: Wolfgang Nagel

Bio:

Abstract:

Hewlett-Packard

Title: HP Product and Roadmap Update: Apollo HPC System Family, Update for HPC and Big Data, DL-Series Update

Speaker: Jimmy Daley, Craig Yamasaki, Mai Nguy, Serge Guex

Bio: Jimmy Daley is the Director of Modular Compute and Storage. He has responsibility for the development of the Apollo 2000 and 4000 as well as all ProLiant Storage solutions. He has held various roles within the HP Server team for the last 17 years.



Serge Guex is working since 1997 for HP in several technical and sales position, In 2012 he moved to the Product Marketing Team in EMEA. Since then he is responsible as a Product Manager for the 1P/2P Rack Servers in EMEA.

Abstract: An overview of the products and technologies in HP's Rack Future portfolio

Hewlett-Packard

Title: BL-Series(Blade Server) Update

Speaker: Jeppe Bjerring

Bio: Been with HP for 9 years, most of that spent in HP Servers. Before joining the HPC & POD team in EMEA, I was responsible for Tier 1 Service Providers and HPC customers in Denmark.



Abstract:

Hewlett-Packard

Title: Superdome X – High-End SMP-System

Speaker: Jerry Huck

Bio: Jerry Huck reports to the Chief Technologist of HP's Server organization working on technical strategy, assessment, and IP portfolio management. His responsibilities include leadership and management of high-level technologists and strategic communication to HP management, HP's field organization, and customers on HP technologies and products. His areas of expertise include a deep understanding of the software/hardware boundary, floating-point arithmetic, virtual memory and virtualization technologies, and a broad understanding of CPU micro architecture, operating systems, compiler optimization, computer security, interconnects, storage arrays, and value-add software layers for platform management.

Over a 30 year career at Hewlett Packard as a systems architect, Huck has worked in intellectual property portfolio management, industry standards development and management, IP out-licensing, patent generation, and assessment. He earned his Ph.D. from Stanford University and holds over 20 patents.

Abstract: This presentation will introduce the recent HP Superdome X SMP architecture which provides mission critical reliability and scalability to both enterprise and High Performance workloads. This system utilizes industry standard Intel Xeon processors into a 16 socket scale-up multiprocessor system with more than 240 cores and 12TB of main memory.

Hewlett-Packard

Title: Moonshot Family Update

Speaker: Gerald Kleyn

Bio: Gerald Kleyn leads the research and development engineering teams that deliver Moonshot server platforms that provide customers with industry leading solutions in Mobile Workspaces, Big Data and Media Processing using energy efficient processing technologies.



Gerald joined Compaq straight out of college as a firmware developer in 1997 and has since rotated into hardware design positions and into various technical leadership roles. His product development experience spans from products for small businesses, Enterprise, High Performance Computing and Service Provider segments.

Gerald enjoys leading teams to deliver cutting edge systems to market that incorporate or invent new technologies that change customer economics through higher performance, lower cost, and lower power consumption. He led development of the first SL multimode systems which have been deployed at numerous service provider and high performance computing installations. These include the Top500 Tsubame 2.0 with the ProLiant SL390s G7 platform and the first FDR Infiniband cluster at Purdue using the ProLiant SL230s Gen8 platform.

Gerald has most recently led the development of HP Project Moonshot, a multi-year project to deliver new levels of performance and energy efficient solutions for the server market.

Gerald holds a BS in Electrical Engineering from Utah State University and an MBA from Tulane University.

Abstract:

PayPal

Title: PayPal Solves Real-time Analytics Problems with HP Moonshot

Speaker: Ryan Quick, Principal Architect, Arno Kolster, Senior Database Architect

Bio: Ryan Quick is an expert at scale-out systems, UNIX kernel design and profiling, and has been recognized for innovation in hardware, application, and object-marshaling. His current efforts bring machine learning, real-time eventing, set-selection, and digital signal processing technologies to bear on predictive analytics in self-healing to command and control systems. Mr Quick holds patents for messaging middleware systems, and is a pioneer in bridging High-Performance Computing technologies with enterprise best-practice infrastructure.



Arno Kolster's main career focus over the past 25 years has been database and operations architecture for a number of different industries including oil and gas, emergency services, finance and e-commerce. An interest in HPC and technical computing came about as a result of exploring solutions to solve real time data analytic problems across distributed systems at web scale. Arno and his colleague, Ryan Quick, received IDC Innovation Excellence Awards at both SuperComputing 2012 and 2104 and he's been invited to speak both domestically and internationally on HPC and its deployment at PayPal.

Abstract:

Come hear PayPal's IT architects, Ryan Quick and Arno Kolster, explain how they created innovative new IT services to support their core business processes by leveraging the capabilities of HP Moonshot and the unique ability to mix different workload-optimized servers into a single chassis, achieving a custom-tuned and highly scalable solution for their specific IT requirements.

Genetic Research Center

Title: Progress in Genetic Research with Advanced HPC

Speaker: Rafael Larrosa

Bio: Rafael Larrosa Jiménez has the B.S. and M.S degrees in Computer Science from the University of Málaga. Has been working as sysadmin at the SCBI (Supercomputing and Bioinformatics) centre at the University of Malaga since 1997, and is Associate Professor in the Computer Architecture Dept. since 2002. Current research interest include parallel file systems and high productivity languages for HPC.



Abstract:

As NGS (Next Generation Sequencing) gets cheaper, and more people gets sequenced, the amount of data that needs to be processed to obtain information from the DNA is bigger. The need to search for similarities between ADN sequences is the big problem, as it takes an exponential time to compare them.

BBVA

Title: SEGA – Strategic Evolution of Grid Architecture

Speaker: Samuel L. Martínez Fernández

Bio: Samuel L. Martinez Fernandez is Head of HPC & Middleware in Corporate & Investment Banking at BBVA. Samuel has more than 7 years experience in HPC because he has been always involved in this kind of projects. Besides, during the last 4 years, he has complemented his knowledge with other middleware tools such as Applications servers or Messaging applications.



His current position makes him to be very worried about systems status, so, monitoring these systems is a priority for him to avoid possible incidents.

Getting all this knowledge together, provides to Samuel a global vision of IT in C&IB BBVA. In this way he can take part in many decision-making where IT is involved.

Abstract:

The main target of SEGA is to create an enterprise grid for BBVA. This includes the use of machines with the operating system in memory as compute nodes or distributing compute nodes in different sites for disaster recovery.

Wellcome Trust Sanger Inst.

Title: HP SL4540 as a scale-out iRODS server

Speaker: Tim Cutts

Bio: Tim Cutts is Head of Scientific Computing at the Wellcome Trust Sanger Institute, responsible for 17,000 cores, 25PB of data, and web services used by Institute scientists in their research. Originally a cell biologist with an interest in mammalian cell cycle checkpoints, he has since his PhD been involved in bioinformatics programming and scientific IT service provision, at the University of Cambridge, then Incyte Genomics, and for the last 13 years at the Sanger Institute.



Abstract:

Rolls-Royce Derby

Title: HPC and Engineering Designs in Rolls-Royce

Speaker: Yoon K Ho

Bio: Yoon is the Rolls-Royce HPC Technical Lead, with technical and operational responsibilities for HPC in the Company. He has been involved in HPC for over 10 years. Before that he was a CFD developer, and also managed the development of CFD technology in RR. Yoon currently serves in the UK e-infrastructure Special Interest Group Steering Committee, and is a member of the HPC 500. Yoon leads the Rolls-Royce global HPC Team, part of the Computational Science and Engineering department. The HPC Team is responsible for the overall global HPC roadmap, ensuring this meets Rolls-Royce simulation requirements.



Abstract: An overview of RR HPC and its role in simulation based engineering designs will be presented, and the challenges for future large scale modelling to HPC will also be discussed

Industrial University of Santander

Title: HPC Knowledge from Oil and Gas to Aerospace Data: Advanced Computing Experiences of the SC3UIS Centre at Bucaramanga, Colombia

Speaker: Carlos J. Barrios

Bio: Director of the High Performance and Scientific Computing Centre at Universidad Industrial de Santander, Bucaramanga, Colombia). General Advisor of the Advanced Computing System for Latin America and Caribbean (SCALAC) and assistant professor of the Informatics and Systems Engineering School of the Universidad Industrial de Santander. Doctor in Computer Science of the Université de Nice-Sophia Antipolis (France), Master in Applied Mathematics and Computer Science of the Université de Grenoble-Alpes (France), Systems Engineer of the Universidad Industrial de Santander (Colombia).



Abstract:

Intel

Title: HPC + Big Data = Opportunity

Speaker: Brent Gorda

Bio: Brent Gorda is the General Manager of the High Performance Data Division at Intel Corp. Prior to joining Intel, Brent co-founded and led Whamcloud, recognized by OpenSFS and EOFS as the primary innovators, developers and providers of technical support for Lustre file system software. Prior launching Whamcloud, Brent held several leadership positions at the Lawrence Livermore National Laboratory where he was responsible for the BlueGene P/Q architecture, as well as some of the large IB-based clusters in use among the DOE laboratories. Brent is the founder of the Student Cluster Competition originally held at Supercomputing 2007, now a worldwide event that showcase the power of parallel computing in the hands of students. Additionally, Brent is a member of the Executive Steering Committee for the SC conference series.



Abstract:

Use of HPC technologies for Big Data applications is taking off. HP and Intel are solving data movement and analytics problems with HPC technologies. Intel's High Performance Data Division (HPDD) developed a Hadoop Adapter for Lustre to run MapReduce, Hive, and Spark jobs on the leading parallel filesystem for the HPC market. The adapter enables BigData jobs to be integrated in an existing job scheduling infrastructure. These software technologies make scale-out storage easier to manage and permit data analytics workloads to be deployed alongside HPC jobs.

Hewlett-Packard

Title: Big Data and HPC: HP's Strategy, Portfolio and Implementation Examples

Speaker: Jimmy Daley

Bio: Jimmy Daley is the Director of Modular Compute and Storage. He has responsibility for the development of the Apollo 2000 and 4000 as well as all ProLiant Storage solutions. He has held various roles within the HP Server team for the last 17 years.



Abstract:


HPC and Big Data continue on a path to intersect as the data sets grow and solutions are needed to manage 10s and 100s of PB for analysis.

NVIDIA

Title: The Accelerated Data Center

Speaker: Marc Hamilton, Vice President, Solutions Architecture and Engineering

Bio: At NVIDIA, the Visual Computing Company, Marc leads the worldwide Solutions Architecture and Engineering team, responsible for working with NVIDIA's customers and partners to deliver the world's best end to end solutions for professional visualization and design, high performance computing, and big data analytics.



Abstract:

With the emergence of GPU-accelerated deep learning, accelerators have crossed the chasm from specialized HPC applications to become a core part of the data center. Because core algorithms like FFT are common to deep learning, HPC, and graphics, NVIDIA's Tesla GPU continues to be the accelerator of choice across all of these application areas. Marc will discuss how key features of NVIDIA's next generation GPU will continue to drive performance and energy efficiency in the accelerated data center.

Mellanox

Title: Interconnect Technology Trends Towards 2020

Speaker: Gilad Shainer

Bio: Gilad Shainer has served as Mellanox's vice president of marketing since March 2013. Previously, Mr. Shainer was Mellanox's vice president of marketing development from March 2012 to March 2013. Mr. Shainer joined Mellanox in 2001 as a design engineer and later served in senior marketing management roles between July 2005 and February 2012. Mr. Shainer holds several patents in the field of high-speed networking and contributed to the PCI-SIG PCI-X and PCIe specifications. Gilad Shainer holds a MSc degree (2001, Cum Laude) and a BSc degree (1998, Cum Laude) in Electrical Engineering from the Technion Institute of Technology in Israel.



Abstract:

Intel

Title: Intel's Scalable System Framework: The Path to Exascale

Speaker: Mark Seager

Bio: Mark K. Seager is an Intel Fellow in the Data Center Group and the chief technology officer for the Technical Computing Ecosystem at Intel Corporation. His work focuses on the development, deployment and use of exascale high-performance computing (HPC) systems with technology based on broadly available HPC systems.



Abstract:

HPC is transforming and addressing the Exascale system challenges requires innovation in multiple areas. Intel's investments in Processor, Fabrics, storage and Software and tools are paving the way to a holistic approach to scalable system solutions. This talk discusses this Scalable System Framework (SSF) and the implications on future systems and applications from traditional HPC to big data and high performance data analytics.

HP Labs

Title: HP's Future High-End-System: "The Maschine" – Towards In-Memory Exascale Computing

Speaker: Paolo Faraboschi

Bio: Paolo Faraboschi is an HP Fellow at HP Labs. He works on The Machine project, researching how we can build better systems around non-volatile memory. His interests are at the intersection of system architecture and software. He previously worked on low-energy servers (HP Project Moonshot), scalable system-level simulation (COTSon), and embedded VLIW cores (Lx/ST200). Paolo is an IEEE Fellow and an active member of the computer architecture community. He holds 26 patents, co-authored over 70 technical publications, and a book. Before joining HP in 1994, he received a Ph.D. in Electrical Engineering from the University of Genoa, Italy.



Abstract:

Different scaling rates in compute and memory have created an imbalance that threatens to decrease the efficiency of tomorrow's supercomputers to an unacceptable point. Moreover, the future progress of DRAM, the technology at the foundation of computing, is expected to slow significantly in the next few years. Today's systems are also challenged to provide adequate memory capacity. We therefore need transformational changes to the way in which we architect, connect, and program the memory subsystem.

The talk will discuss memory technology trends and the implications of those trends to architecture, software, and programming. Starting from emerging non-volatile devices, the talk covers how we can flatten and re-architect the memory hierarchy and how this can help the next generation of supercomputer applications. It will also present "The Maschine", a research project at HP Labs that aims at advancing, implementing and testing these new memory system technologies.

Tokyo Institute of Technology

Title: Programming Paradigms and Algorithms for the Exascale World

Speaker: Satoshi Matsuoka

Bio:

Abstract:

Hewlett-Packard and KIT

Title: Conference Closing Session: Evaluation of Questionnaires & General Q&A

Speaker: Alain Andreoli

Bio:



Alain Andreoli is the leader of the HP Servers Business Unit where he drives technological and go-to-market strategies for Compute to generate even greater business value for HP customers and partners. HP Servers is a \$13+ billion business, including HP Cloudline, built for scale; HP Apollo family of high performance computing; HP Moonshot, the world's first software defined web server; HP ProLiant servers, the world's number one server brand; HP BladeSystem, the world's number one bladed infrastructure; and HP Integrity for mission-critical computing.

In this role, Alain is responsible for serving customers of all sizes and delivering compute solutions for Social, Mobile, Big Data and Cloud around the globe. As the world leader in server and solutions, HP powers nearly every aspect of daily life, from most of the world's stock exchanges, ATM transactions, 911 systems and across enterprises. Alain's core competency is to lead fast-paced, complex and strategic business transformations. He has executed a variety of turnarounds in thirty years of business experience around the globe in semiconductors, computing, storage, networking, software, cloud, services and multimedia. Alain joined HP in 2013 from Grass Valley, the leading global live video solutions provider, where he was President, CEO and Co-Investor based in San Francisco. Prior to this role, he was the President of Sun Microsystems Europe based in London. Alain spent fourteen years at Texas Instruments with various roles including Chairman of the global industrial segment, was the Corporate VP/GM of StorageTek (Oracle) for international operations and global services, President and COO of Verio (NTT), EVP of McData (Brocade), CEO of Xiotech, and Chairman and CEO of Hubwoo. Alain earned a Master's in Management from SKEMA Business School in France and participated in the MBA program at Concordia University in Canada. Alain is a dual American-French citizen, is married with three children, and lives between San Francisco and Houston.

Saturday, July 11th, 2015

Tutorial A1: Intel Technology Tutorial

Intel

Title: The Future of Technical Computing: Intel's HPC Scalable System Framework

Speaker: Thor Sewell, Joe Yaworski, Brent Gorda

Bio: **Thor Sewell** is the Director of Technical Computing Product Marketing for Intel's Data Center Group. His team's responsibilities include Intel Xeon Phi product family and Intel Xeon Processor for the technical computing segment. Thor is based in Portland, Oregon.



Joe Yaworski is Director of Fabric Product Marketing for Intel's Technical Computing Group. He is responsible for all of Intel's marketing activities for Intel high performance fabric programs. This includes positioning and promotion of Intel's True Scale Fabric solutions as well as the plans to bring the next generation Intel fabrics to market. Also as part of his responsibilities, he directs the Intel Fabric Products Center; which is used for performance profiling and customer benchmarking. Joe came to Intel as part of Intel's acquisition of the QLogic InfiniBand program where he was Director of Marketing for QLogic's HPC programs.



Brent Gorda is the General Manager of the High Performance Data Division at Intel Corp. Prior to joining Intel, Brent co-founded and led Whamcloud, recognized by OpenSFS and EOFS as the primary innovators, developers and providers of technical support for Lustre file system software. Prior launching Whamcloud, Brent held several leadership positions at the Lawrence Livermore National Laboratory where he was responsible for the BlueGene P/Q architecture, as well as some of the large IB-based clusters in use among the DOE laboratories. Brent is the founder of the Student Cluster Competition originally held at Supercomputing 2007, now a worldwide event that showcase the power of parallel computing in the hands of students. Additionally, Brent is a member of the Executive Steering Committee for the SC conference series.



High Performance Computing (HPC) has reached an inflection point with the convergence of traditional HPC and the emerging world of big data analytics. So why not join our breakout session to learn about how Intel's HPC Scalable System Framework provides systems based approach to a balanced platform for the HPC platforms today and in the future? We will cover the Intel® Xeon Phi™ Product family to include the next generation Knights Landing as well as Intel® Xeon® Processors, Intel® Omni-Path Architecture, and Intel® Lustre* solutions. Intel Confidential.

Abstract:

Tutorial A2: GPU Technology Tutorial

NVIDIA

Title: Tesla GPU Accelerator Roadmap

Speaker: Timothy Lanfear

Bio: Timothy Lanfear manages the European solution architecture and engineering team in NVIDIA's Professional Solutions Group. He has twenty years' experience in HPC, starting as a computational scientist in British Aerospace's corporate research centre, and then moving to technical pre-sales roles with Hitachi, ClearSpeed, and most recently NVIDIA. He has a degree in Electrical Engineering and a PhD for research in the field of graph theory, both from Imperial College London. He is a Chartered Engineer and Member of the Institution of Engineering and Technology.



The Pascal generation of NVIDIA GPUs, coming in 2016, will offer both significantly higher performance and simplifications in ease of use. NVLink, a new high-speed data path, will bypass the bandwidth restrictions of the PCIe bus; and once the data is on the GPU high bandwidth stacked memory will significantly boost application performance. The new unified memory feature simplifies the application developer's task of managing data migration between the host and GPU memories.

Abstract:

[Tutorial A3: AMD Technology Tutorial](#)

AMD

Title: **AMD: HW and SW innovation for HPC**

Speaker: **Karl Freund**

Bio: Karl Freund has been an executive in the IT hardware and software business for over 35 years, both at system vendors such as HP, Cray, and IBM, at Calxeda, the ARM Server Startup, and now at AMD as VP of High Performance Computing. Mr. Freund has been widely quoted in hundreds of articles, and has been a frequent speaker at technology and investment conferences. At AMD, he leads the company's strategy and execution of acceleration technologies that are enabling AMD to re-emerge as a leader in HPC.



Mr. Freund holds a bachelors degree from Texas A&M University in Applied Mathematics, and a Masters of Computer Science from the University of North Carolina at Chapel Hill.

Abstract:

When AMD earned the #1 spot on the Green 500 last year, most people thought they had traveled back in time, a time when AMD was featured in many "Top500" systems based on Opteron. But the reality is that AMD's FirePro S9150 is the most power efficient accelerator in the industry today, and highlights the company's ambitions to serve this important market and customers. This talk will discuss both the AMD accelerator hardware, now supported on multiple HP Servers, as well as the fascinating Software roadmap that will enable more applications to take advantage of acceleration without resorting to difficult software programming models. C, anyone?

[Tutorial A4: Accelerators Experiences](#)

Kalray

Title: **Cloud Acceleration MPA 2nd Generation Kalray Processors**

Speaker: **Jean-Pierre Demange**, Vice President Sales and Marketing

Bio: Jean Pierre is VP Sales and Marketing at Kalray. He has more than 30 years of business experience in the semiconductor industry related to Fixed and Wireless Telecommunication at different Management positions in Sales, Marketing and Product Business Unit that he held at Texas Instruments. Products that he managed range from Processors, Digital Signal Processors (DSP), ASIC and dedicated Telecommunications devices. Jean Pierre has an electronic engineer degree from the University of Montpellier France.



Abstract:

MPPA 2nd generation Kalray Processor called Bostan will sample very soon and will be in Production by 4Q15. Andey, the 1st generation is already used in the PCIe Turbocard2 to accelerate some workloads in HPC like cryptography, RTM for oil & gas and MonteCarlo for finances. FPU and memory bandwidth has been increased on Bostan which will boost considerably the performance of the new Compute acceleration Turbocard3 available in 4Q15. Kalray will also accelerate Cloud Networking with Bostan as a programmable Ethernet dispatcher and 128 crypto accelerators has been integrated on chip. Thanks to this new Ethernet interface, deterministic architecture and its 256 cores Bostan will be in the leadership position as it will allow customers to offload standard Converged Ethernet protocols, security and storage protocols (RoCE, IPsec, iSCSI...) but also to implement specialized HPC processing in C/C++ or even Fortran directly on the NIC itself for latency-sensitive workloads. Kalray will introduce by end of the year two Network Interface Cards, the Kalray Open NIC-80 (KONIC-80) half-length running at 2X40Gbits and the Kalray Open NIC-40 (KONIC-40) low profile running at 1X40Gbits.

University of the Federal Armed Forces

Title: Parallel Processing and HPC in Numerical Simulations with GPUs

Speaker: Arash Ramezani

Bio: Arash Ramezani currently works as a research assistant in the Department of Mechanical Engineering at the Helmut-Schmidt-University in Hamburg. He has studied Applied Mathematics at the University of Bremen and the University of Queensland in Australia and received his Diploma degree in 2010. In 2015 he received his doctor's degree in engineering science. His research interests include modeling, simulation and visualization of ballistic problems.



Abstract:

The objective of this work is to develop and improve the modern armor used in the security sector. Instead of running expensive ballistic tests, numerical simulations should identify the vulnerabilities of items and structures. To develop better, smarter constructions requires analyzing a wider range of parameters. However, there is a simple rule of thumb: the more design iterations that can be simulated, the more optimized is the final product. As a result, a new high-performance computing (HPC) solution has to dramatically reduce overall engineering simulation time. All possible methodologies will be discussed with an application-oriented evaluation.

NVIDIA

Title: Deep Learning with GPUs for Service Providers

Speaker: Timothy Lanfear

Bio: Timothy Lanfear manages the European solution architecture and engineering team in NVIDIA's Professional Solutions Group. He has twenty years' experience in HPC, starting as a computational scientist in British Aerospace's corporate research centre, and then moving to technical pre-sales roles with Hitachi, ClearSpeed, and most recently NVIDIA. He has a degree in Electrical Engineering and a PhD for research in the field of graph theory, both from Imperial College London. He is a Chartered Engineer and Member of the Institution of Engineering and Technology.



Abstract:

Making sense of vast amounts of data is vital to the business of service providers. The GPU architecture has proven to be ideal for running deep learning algorithms that can automatically understand data like images, speech and text. NVIDIA's deep learning platform offers the tools needed by researchers and developers in this field.

Tutorial A5: Accelerator S/W Environments

Nallatech

Title: FPGA Acceleration using OpenCL

Speaker: Craig Petrie

Bio: Business Development Manager with twelve years of product development and technical marketing experience specializing in accelerated processing solutions featuring Altera and Xilinx FPGA technology.



Abstract:

For almost a decade now, FPGAs have demonstrated tremendous potential as accelerators for Scientific and High Performance Computing markets. Programmability and cost have been two of the historic inhibitors to adoption, however these obstacles have now being overcome with market changing implications. This presentation summarizes the latest developments that position FPGAs as viable energy-efficient accelerators delivering new levels of application price-performance.

Rogue Wave Software

Title: **Using TotalView and other Rogue Wave tools to help leverage the performance of accelerated clusters**

Speaker: **Nikolay Piskun**

Bio: Nikolay Piskun, Ph.D., Director of Continuing Engineering, has been with TotalView Technologies for almost ten years. He has a Doctor of Philosophy in Physics from Ohio State University and is specializing in parallel computing. Nikolay works closely with the product management team to set the direction for development of the entire family of TotalView Technologies products. He has provided tutorials, workshops, and presentations at numerous technical conferences, including SuperComputing and International Supercomputing for the past several years.



Abstract: Accelerators and coprocessors are providing a significant performance boost for scientific codes on many of the most recently installed HPC systems. These performance gains come at the cost of increased program complexity with program execution that is essentially heterogeneous and where, in some cases, part of the program is running in a restricted environment without the usual capabilities afforded by an operating system. TotalView provides a comprehensive set of capabilities for debugging programs running on heterogeneous and many-core architectures. TotalView supports CUDA and OpenACC on NVIDIA accelerated systems and the Intel tool chain including OpenMP, MPI, and LEO on systems with Intel Xeon Phi Coprocessors. This talk will focus on these features and provide a brief roadmap of future development.

Altera

Title: **Altera Product Update**

Speaker: **Mike Strickland**, Director, Computer & Storage Architecture

Bio: Mike Strickland has more than twenty years of computer, networking and storage experience with companies such as Hewlett Packard, Silverback Systems, and Altera. He currently is leading the computer and storage business unit strategic direction, and is also leading the FPGA High Performance Computing vision across Altera. Previously Strickland has led the development and launch of numerous products including networking, storage management, TCP/IP Offload and iSCSI. He holds a B.S. degree in electrical engineering from Brown University and a M.S. degree in management from the Sloan School of Management at M.I.T.



Abstract: Microsoft has announced that they will accelerate Bing search, machine learning, and networking with FPGAs. Intel announced that they plan to provide Xeon and FPGA dice in a single Xeon package to accelerate servers. There are increasing opportunities in the data center for FPGA algorithm, networking, and data access acceleration. A key enabler to FPGA usage in the data center is OpenCL, but advances in FPGA hardware architecture are also opening up new opportunities. This presentation will highlight the role and the impact of FPGAs in the data center.

Allinea

Title: **Reducing Application Energy Footprint with Accelerators and Allinea Tools**

Speaker: **Patrick Wohlschlegel**

Bio: Patrick graduated with an MSc from Bordeaux Graduate School of Engineering (ENSEIRB), specializing in parallelism and distributed systems. After being trained in Japan at the University of Kumamoto, he went to IBM in Paris as a pre-sales benchmarker. Through the experience he gained, he was then given the opportunity to create an HPC activity within a vendor reseller, Ovesys. As he developed the relationships between Ovesys and key HPC actors to design, build and deploy HPC centers, he naturally discovered Allinea. Today, Patrick is based in Warwick (UK) and works as a Technical Services Manager - sharing his expertise in HPC software tools to users all around the world.



Abstract: Technologies used within the HPC market are evolving and over the last few years, we have witnessed an important shift of applications from CPU to accelerators (such as Nvidia GPUs or Intel Xeon Phi). One of the core reasons for this change is the promise of a very high Flops/Watt ratio, helping the HPC community reach a high level of performance whilst minimizing energy costs. Nevertheless - such migrations can be challenging and often require code modernization. To help developers adopt new technologies quicker, Allinea has expanded the capabilities of Allinea Performance Reports and Allinea Forge to address these matters.

Tutorial B1: HPC Storage System Components

Hewlett-Packard

Title: **HP Storage Server Portfolio: System Components and Subsystems for HPC & Big Data – A Detailed Assessment and Application Examples**

Speaker: **Jimmy Daley**

Bio: Jimmy Daley is the Director of Modular Compute and Storage. He has responsibility for the development of the Apollo 2000 and 4000 as well as all ProLiant Storage solutions. He has held various roles within the HP Server team for the last 17 years.



Abstract: No two Big Data solutions are alike. How do you make sure yours is tailored for your problems? We will discuss some of the configuration options available and where some of these alternatives were best fit to solve real world problems.

Tutorial B2: Focus on Future Development, Deployment and Support of LUSTRE®

Intel & Hewlett-Packard

Title: **Creating a scale out storage using commodity HP storage**

Speaker: **Gabriele Paciucci, Intel & Sandor Bihary, HP**

Bio: **Ing. Gabriele Paciucci** is a solution architect in the High Performance Data Division at Intel. In this role, Ing. Paciucci provides technical consultations to partners and customers and evangelizes the Lustre technology worldwide. Gabriele joined Intel in 2013. Previously, Gabriele was a senior software engineer specialized in HPC and Cloud solution based on Open Source Software. He has architected a number of high performance computing storage solutions based on Lustre on a variety of hardware platforms since 2006. Gabriele is involved in several Open Source projects and has promoted the adoption of Open Source software and Linux in the Enterprise since 2000 when he worked as software engineer at Red Hat. Ing. Paciucci received his Master Degree in Chemical Engineering from “Università degli Studi di Roma La Sapienza” in 1999.



Sándor Bihary is a technical consultant at HP. Sándor joined HP in 2010. Sándor provides technical consultations to customers, perform and lead installations and implementations. He is focused on HPC, Enterprise storages and Virtualization. He earned a master's degree in Electrical Engineering from Budapest University of Technology and Economics in 2010.

Abstract: Intel and HP designed together a cost effective solution for NIIFI (National ICT Infrastructure Development Institute) in Hungary using Intel Enterprise Edition for Lustre. NIIFI plans to get into the Top500 list soon and requested a scale out solution using HP commodity storage to maximize the return of the investment. The presentation covers all the phases of the project from the design to the deliver and include the support provided by Intel to the local HP team

Seagate

Title: **Next-Gen Seagate ClusterStor: Enterprise-Ready Parallel Storage with Leading Energy Efficiency & Performance**

Speaker: **Torben Kling Petersen**, Principal Engineer, Seagate Technologies

Bio:



Torben Kling Petersen has worked with high performance computing in one form or another since 1994 and is currently working as the Principal Engineer in Seagate HPC Solutions Group. Over the years, both in Sun Microsystems and subsequently in Xyratex, he has worked in a number of capacities such as lead architect for enterprise datacenter infrastructure for several telecommunication OEMs, technical lead for IPTV and product specialist for high-end visualization, and global architect for the Lustre BdM team as well as a Cloud Computing infrastructure architect to mention a few. Over the years, Torben has worked on many large HPC systems around the world including Sandia Red Sky in the US, ANU/BOM in Australia, CHPC in South Africa, ETH Zürich and many others in Europe. In addition to compute and storage systems, Torben has also worked on data center design including the new CSCS datacenter in Lugano, Switzerland.

Since joining Seagate more than 4 years ago, Torben has helped bring a number of products to market as well as assist on a number of large ClusterStor/Sonexion wins such as CSCS, ECMWF, EPCC Archer, British MetOffice, DKRZ, CHPC to mention a few.

A great many systems currently under deployment or in planning stages appear to be putting ever increasing requirements on storage density, capacity, manageability and power consumption. While performance is still of utmost concern for bulk of Top500 systems, reliability, availability and supportability are the new critical keywords to the industry. In addition, enterprise features involving data integrity, data management and automated archiving are raising the bar for companies delivering fully integrated and engineered solutions. This talk will touch on some of the requirements recently put forward by the largest storage procurements in HPC in 2014 and what Seagate has accomplished in designing and delivering a solution that maximizes density and performance while fitting into a restrictive power envelopes. Seagate will describe how we can ship a 45PB+ usable storage volume delivering around 500 GB/s using less than 200 kW in power based on Lustre 2.5.x.

Abstract:

ParTec / EOFS

Title: **Status and Future of "lustre.org"**

Speaker: **Hugo Falter**

Bio: Hugo Falter is Co-founder and Chief Operating Officer of ParTec Cluster Competence Center GmbH, a spin-off from the Computer Science Department of the University of Karlsruhe in 1999. Mr. Falter studied law in Regensburg and Munich. Together with a Munich law firm he specialized in bringing innovative technology companies to market. In the Munich law firm Frohwitter, Hugo Falter is responsible for the firm's subsidiary ParTec Cluster Competence Center GmbH.



EOFS and OpenSFS are working together to strengthen the Lustre community.

Abstract: A significant contribution to this goal was the transfer of the "Lustre.org" - website to EOFS and OpenSFS, who own and maintain "Lustre.org" jointly.

The talk will show how this Lustre community achievement was set up.

Tutorial B3/4: HPC and Big Data Storage – and Innovative Parallel File-Systems

DDN

Title: **Leveraging the HP/DDN Technology Partnership for Competitive Advantage**

Speaker: **Molly Rector**, Chief Marketing Officer

Bio:



With 15 years of experience working in the HPC, Media and Entertainment, and Enterprise IT industries running global marketing programs, Molly Rector serves as DDN's Chief Marketing Officer responsible for product management and worldwide marketing. Rector's role includes providing customer and market input into the company's product roadmap, raising the Corporate brand visibility outside traditional markets, expanding the partner ecosystem and driving the end-to-end customer experience from definition to delivery. Rector is a founding member and currently serves as Chairman of the Board for the Active Archive Alliance. She is also the Storage Networking Industry Association's (SNIA) Vice Chairman of the Board and the Analytics and Big Data committee Vice Chairman. Prior to joining DDN, Rector was responsible for product management and worldwide marketing as CMO at Spectra Logic. During her tenure at Spectra Logic, the company grew revenues consistently by double digits year-over-year, while also maintaining profitability. Rector holds certifications as CommVault Certified System Administrator; Veritas Certified Data Protection Administrator; and Oracle Certified Enterprise DBA: Backup and Recovery. She earned a Bachelor's of Science degree in biology and chemistry.

Abstract:

In this session, we will deep dive into the 2015/2016 roadmap for DDN technology in HPC block, file, cache and burst buffer technologies. We will cover both the DDN hardware and software offerings. On the hardware side, we will reveal the latest benchmarks for our next generation SFA storage platform which offers the fastest mixed IO, IOPS and throughput of any storage platform currently available in the space. And on the software side of the business, a deep dive into SFX flash cache and IME burst buffer technologies will detail how best to optimize the usage of your HPC hardware platforms. Join this session to not just hear about our roadmap but to meet representatives from the DDN Paris Development Center and to hear about the collaborative work being done in Grenoble for both benchmarks and unique offerings available only through the HP/DDN partnership.

Intel

Title: **Beyond Posix: An Update on the DAOS**

Speaker: **Brent Gorda**

Bio:



Brent Gorda is the General Manager of the High Performance Data Division at Intel Corp. Prior to joining Intel, Brent co-founded and led Whamcloud, recognized by OpenSFS and EOFS as the primary innovators, developers and providers of technical support for Lustre file system software. Prior launching Whamcloud, Brent held several leadership positions at the Lawrence Livermore National Laboratory where he was responsible for the BlueGene P/Q architecture, as well as some of the large IB-based clusters in use among the DOE laboratories. Brent is the founder of the Student Cluster Competition originally held at Supercomputing 2007, now a worldwide event that showcase the power of parallel computing in the hands of students. Additionally, Brent is a member of the Executive Steering Committee for the SC conference series.

Abstract:

The Distributed Application Object Store (DAOS) interface is the result of a US DOE project called Fast Forward. The intent of the project was to figure out how applications in the massively parallel world might best interface with storage in the future. Some of the key work out of this have been implementation of the concept of a "burst buffer" with products available today.

Panasas

Title: **ActiveStor: Hybrid Scale-Out NAS that Increases in Reliability with Scale**

Speaker: **Geoffrey Noer**, Vice President, Product Management

Bio:



Geoffrey Noer is Vice President of Product Management at Panasas. In this position, he is responsible for all outbound and inbound product management and product marketing, including product positioning, market requirements for future Panasas products, and go-to-market strategies. Previously, he led the product management team at SGI for its server product lines. Geoffrey joined SGI through its acquisition by Rackable Systems where he was VP of Product Management for all products and sat on the Green Grid board of directors. Earlier, he worked for Zambeel, Red Hat, and Cygnus Solutions. Geoffrey holds a BA in Computer Science from Swarthmore College.

Panasas hybrid scale-out NAS appliances simplify storage performance and data management, driving customer innovation, value, and competitive advantage. In this talk, Panasas will share the latest on its ActiveStor hardware platform and PanFS storage operating system releases. Delivering triple parity data protection based on erasure codes instead of hardware RAID, PanFS 6.0 represents the most significant advance in several years for Panasas. It revolutionizes reliability and availability by enabling data protection to actually increase with scale instead of decreasing as is expected with competing file systems.

Abstract:

Scality

Title: **Scale-out Object Based Storage on the Exascale Road**

Speaker: **Bradley King**

Bio:



Bradley King is one of the co-founders and the chief architect at Scality. As chief architect, he is responsible for the design of the largest storage systems Scality is deploying around the world. These include multi-petabyte, multi-site systems, with hundreds of servers. Brad has had a multifaceted career in fluid dynamics research and CFD as well as working with fixed and mobile operators throughout Europe, selling messaging solutions, WAP gateways, MMSCs, embedded handset software, and location-based solutions. Brad holds a PhD degree in Naval Architecture and Marine Engineering from the University of Michigan.

The exascale challenge is uniting many researchers and vendors around the world to look into the future and seek out the path to yet another thousand times more simulation capability. There is a general consensus that today's storage models will not be able to meet this need and object based storage models are among the strongest candidates for the next generation in storage. This presentation will present key ideas around the object based approach for exa-scale storage and its pertinence to meeting the challenge. Additionally, examples of today's Scality platforms in various fields of research and development on the exascale road will be presented.

Abstract:

FhG/ITWM (Fraunhofer Society)

Title: BeeOND - a BeeGFS based product for bursting I/O

Speaker: Franz-Joseph Pfreundt

Bio: Dr. Franz-Josef Pfreundt studied Mathematics, Physics and Computer Science resulting in a Diploma in Mathematics and a Ph.D degree in Mathematical Physics (1986).



From 1986-1995 he had a permanent position at the University of Kaiserslautern as Head of the Research Group for Industrial Mathematics. In 1995 he was cofounder of the Fraunhofer Institute for Industrial Mathematics – ITWM. He founded the departments: "Flow in complex structures" and "Models and algorithms in image analysis".

Since 1999 he is Division Director at Fraunhofer ITWM and Head of the "Competence Center for HPC and Visualization". In 2001 the prestigious Fraunhofer Research Prize was awarded to Franz-Josef Pfreundt, Konrad Steiner and their research group for their work on microstructure simulation. He is PI for a variety of parallel computing research projects including the development of new parallel programming frameworks. He initiated the Fraunhofer Resource Grid in 2001 and the Fraunhofer Cloud Alliance. The developments in the area of visualisation and implementations on IBM Cell Processor gained the Fraunhofer Research Price in 2005 and the IBM faculty award in 2006. His main research focuses since 2005 are parallel filesystems (BeeGFS), new parallel programming approaches (GPI) and Big Data frameworks (GPI-Space) as well as the use of these technologies for the simulation challenges in oil&gas exploration.

Abstract:

The BeeGFS parallel file system has seen some great adoption over the last 12 months from Academia as well as from industry. The talk will cover the main two of the most interesting topics. The high availability build in into BeeGFS without using external software. BeeOND: BeeGFS on demand and how it today solves the problems related with bursting I/O behavior.

Max Planck Institute for Metabolism Research

Title: BeeGFS in a scientific environment: the Good, the Bad and the Ugly

Speaker: Stefan Vollmar

Bio: Dr. Stefan Vollmar is the Head of IT at the Max Planck Institute for Metabolism Research, Germany. He is a physicist and has a strong background in scientific software development. In particular, his group focuses on the visualization and analysis of tomography data, also on tools for good scientific practise (workflow documentation).



Abstract:

Last year, we have introduced BeeGFS at our institute. It is running on several HP SL4540 servers and used for scientific computing on our HP c7000 blade servers (InfiniBand networking) - despite the title: experiences are mostly good and better.

Tutorial C1/C2: Interconnect Technologies, Shared Memory Environments and Related Tools

Mellanox

Title: Interconnecting HPC Systems at 100Gb/s speed

Speaker: Gilad Shainer

Bio: Gilad Shainer has served as Mellanox's vice president of marketing since March 2013. Previously, Mr. Shainer was Mellanox's vice president of marketing development from March 2012 to March 2013. Mr. Shainer joined Mellanox in 2001 as a design engineer and later served in senior marketing management roles between July 2005 and February 2012. Mr. Shainer holds several patents in the field of high-speed networking and contributed to the PCI-SIG PCI-X and PCIe specifications. Gilad Shainer holds a MSc degree (2001, Cum Laude) and a BSc degree (1998, Cum Laude) in Electrical Engineering from the Technion Institute of Technology in Israel.



Abstract:

The exponential growth in data and the ever growing demand for higher performance to serve the requirements of the leading scientific applications, drive the need for Petascale system and beyond and the ability to connect tens-of-thousands of compute and co-processor nodes in a very fast and efficient way. The interconnect has become the enabler of data and the enabler of efficient simulations. Beyond throughput and latency, the data center interconnect needs be able to offload the processing units from the communications work in order to deliver the desired efficiency and scalability. The presentation will cover the latest technology from Mellanox, including the end-to-end 100Gb/s interconnect solutions, the HPC-X software package that provides a complete solution for MPI and PGAS/SHMEM/UPC environments with smart offloading techniques, and a roadmap for the next generation InfiniBand speeds.

Mellanox

Title: Enhancing Applications for Next Generation Performance

Speaker: Gilad Shainer

Bio: Gilad Shainer has served as Mellanox's vice president of marketing since March 2013. Previously, Mr. Shainer was Mellanox's vice president of marketing development from March 2012 to March 2013. Mr. Shainer joined Mellanox in 2001 as a design engineer and later served in senior marketing management roles between July 2005 and February 2012. Mr. Shainer holds several patents in the field of high-speed networking and contributed to the PCI-SIG PCI-X and PCIe specifications. Gilad Shainer holds a MSc degree (2001, Cum Laude) and a BSc degree (1998, Cum Laude) in Electrical Engineering from the Technion Institute of Technology in Israel.



Abstract:

This session will explore the network capabilities from RDMA to GPUDirect and hardware-based aggregation protocols, and how these technologies continue to evolve to meet the needs for future applications and systems. The session will include a report from the newly formed Open Fabric OFV group as well as results from collaborative work on creating the most advanced software API for HPC applications

Intel

Title: **The Next-Generation Fabric: New Details on the Intel® Omni-Path Architecture**

Speaker: **Joe Yaworski**

Bio:



Joe Yaworski is Director of Fabric Product Marketing for Intel's Technical Computing Group. He is responsible for all of Intel's marketing activities for Intel high performance fabric programs. This includes positioning and promotion of Intel's True Scale Fabric solutions as well as the plans to bring the next generation Intel fabrics to market. Also as part of his responsibilities, he directs the Intel Fabric Products Center; which is used for performance profiling and customer benchmarking. Joe came to Intel as part of Intel's acquisition of the QLogic InfiniBand program where he was Director of Marketing for QLogic's HPC programs.

Abstract:

This session discusses new details on the Intel® Omni-Path Architecture, Intel's next-generation fabric product line, which is designed around industry-leading technologies developed as a result of Intel's multi-year fabric development program. It will cover new additional product information not previously disclosed. It will review how Intel Omni-Path Architecture will deliver significant enhancements and optimization for HPC at both the host and fabric levels, thereby providing significant benefit to HPC applications over currently available interconnect technologies.

EXTOLL

Title: **EXTOLL – Measurement with an Extreme Low-Latency Interconnect on HP Servers**

Speaker: **Ulrich Brüning, Mondrian Nüssle**

Bio:



Prof. Brüning is one of the co-founders of EXTOLL. He is a renowned researcher in the area of computer architecture with a long track record of designing parallel hardware architectures and interconnection networks. Prof. Brüning is a professor of computer architecture at the University of Heidelberg, and co-founded and was the CEO of a company that offered several unique sequential and parallel UNIX computing systems. He has worked on numerous parallel and high-performance machines and architectures. He designed the hardware for the function-parallel machine „Starlet“ in

1983 and the node architecture for the SUPRENUM parallel machine (an initiative to build a German supercomputer from

1985 to 1990). In 1993, he architected the parallel MANNA computer based on the Intel i860 processor and custom ASICs.

Until 1995, he also took part in the development of the massively parallel machine PowerMANNA based on the Motorola PowerPC620. In 2003, he presented the HPC interconnection network ATOLL, which was a precursor to EXTOLL. Within EXTOLL he sets the research strategy and direction and manages collaboration with academia.



Dr. Nüssle is one of the co-founders of EXTOLL. He has more than 12 years of experience in research and development. In these years he participated in a number of research and industrial projects in the areas of computer architecture, low-level software, and RTL hardware design. Dr. Nüssle has worked on the EXTOLL project from the beginning and helped design and implement both the hardware and software. As the leader of the software and hardware engineering teams, he defines the technical vision for EXTOLL.

Abstract:

The new EXTOLL Tourmalet interconnect for HPC offers outstanding performance in terms of several of the metrics important for HPC computing.

Actual measurements on a HP platform will be presented, as well as some further exciting features of this new interconnect for HPC.

ScaleMP

Title: Experience and Future Development with Very Large Memory Solutions for HPC and Big Data

Speaker: Benzi Galili

Bio: Benzi joined ScaleMP in 2005, and is responsible for ScaleMP's operations and sales worldwide.



Previously VP of Services for ScaleMP, Benzi brings more than 20 years of experience in managing technology development and services, product and account management, and in product architecture and design geared towards addressing business needs of enterprises and end-users.

Prior to joining ScaleMP, Benzi served as VP of Products at Participate Systems (Acquired by OutStart Inc, now Kenexa - Division of IBM), VP R&D, and VP Products and Services at Kamoon Inc, and also as CIO for the ERGO Consulting Group. Throughout his career he has managed diverse engineering projects, products, sales teams, and business units, delivering key products and services to Global 2000 companies, governments, and leading academic and research institutes. Benzi holds a B.Sc. in industrial engineering from Tel Aviv University (major in information systems), where he graduated with honors.

Abstract:

Tutorial D1/D2: Moonshot Technology

Hewlett-Packard

Title: Moonshot Update: New Solutions, Performance Measurements and Roadmap

Speaker: Gerald Kleyn

Bio: **Gerald Kleyn** leads the research and development engineering teams that deliver Moonshot server platforms that provide customers with industry leading solutions in Mobile Workspaces, Big Data and Media Processing using energy efficient processing technologies.



Gerald joined Compaq straight out of college as a firmware developer in 1997 and has since rotated into hardware design positions and into various technical leadership roles. His product development experience spans from products for small businesses, Enterprise, High Performance Computing and Service Provider segments.

Gerald enjoys leading teams to deliver cutting edge systems to market that incorporate or invent new technologies that change customer economics through higher performance, lower cost, and lower power consumption. He led development of the first SL multimode systems which have been deployed at numerous service provider and high performance computing installations. These include the Top500 Tsubame 2.0 with the ProLiant SL390s G7 platform and the first FDR Infiniband cluster at Purdue using the ProLiant SL230s Gen8 platform.

Gerald has most recently led the development of HP Project Moonshot, a multi-year project to deliver new levels of performance and energy efficient solutions for the server market.

Gerald holds a BS in Electrical Engineering from Utah State University and an MBA from Tulane University.

Abstract:

Texas Instruments

Title: **Graphs and Deep Learning on Moonshot with m800 Cartridges**

Speaker: **Arthur Redfern**

Bio: Arthur J. Redfern received a B.S. in 1995 from the University of Virginia and a M.S. and Ph.D. in 1996 and 1999, respectively, from the Georgia Institute of Technology, all in electrical engineering. Following his thesis work on nonlinear systems modeled by the Volterra series, Arthur joined the Embedded Processing Systems Labs at Texas Instruments where he currently manages the High Performance Computing Lab. His activities at TI have spanned the areas of high performance computing (machine learning, graphs and signal processing), intersecting analog and signal processing methods (ADCs, amplifiers, DACs, design optimization, speakers and touch screens) and communication system design (DSL, DTV and SerDes). He has over 20 papers published in refereed conferences and journals and has been granted over 20 US patents.



Abstract: This presentation focuses on 2 libraries developed for HP Moonshot with m800 cartridges: the TI Graph Library and the TI Machine Learning Library. Library design, performance and example applications are described that build on the strengths of HP Moonshot's networking and acceleration capabilities with TI66AK2H12 SoCs. Examples include the Graph500 benchmark and deep learning based single image classification and multiple object detection.

Justus-Liebig-University, Giessen

Title: **Evaluation of Moonshot Cartridges M300 and M710 for bioinformatics tasks**

Speaker: **Burkhard Linke**

Bio:

- Born 1976 in Herford, Germany
- Studied computer science in Bielefeld, Germany; finished with diploma in 12/2002
- Member of scientific staff in Bielefeld since 01/2003 in the field of bioinformatics
- Finished PhD in 02/2013
- Changed to Giessen in 08/2013 in the course of funding of a new professorship
- Half time system administrator, half time bioinformatician



Abstract: The field of bioinformatics presents a number of different challenges to compute resources. One of the best known application is 'BLAST', a tool for searching homologue sequences in large sequence databases.

We have evaluated HP Moonshot cartridges as an extension to our existing compute cluster based on HP SL230s with respect to BLAST and similar applications on a real life dataset.

Hewlett-Packard

Title: **FPGA Acceleration on Moonshot**

Speaker: **Mitch Wright**, Master Engineer HP-Server Chief Technology Office Staff

Bio: Mitch Wright is a staff engineer in the HP-Server Chief Technology Office, in this role Wright evaluates and champions new technologies for the HP Server product organization. Mitch has been with HP/Compaq Industry Standard Servers for 18 years and has contributed to notable products and technology development, Proliant 8000 first 8-socket x86 server, DL585G1 first Opteron, GPGPUs, FusionIO, Moonshot and now FPGA Acceleration.



Mitch has 8 patents issued and multiple pending. Prior to HP/CPQ, Mitch worked for 11 years designing instrumentation for High Energy Physics.

Abstract: Recent published Search and Deep Learning papers and the announced 14nm/16nm devices have FPGAs all the buzz.

We will review HP FPGA investigations on the Moonshot platform.

[Tutorial E1/2: Workflow Management & Clouds](#)

Hewlett-Packard

Title: HP's HPC-Cloud Strategy and Execution

Speaker: Jean-Luc Assor

Bio:

Abstract:

The UberCloud

Title: Containerization of HPC Applications in the Cloud

Speaker: Wolfgang Gentzsch

Bio: Wolfgang Gentzsch is president and co-founder of the UberCloud Community and Marketplace for engineers and scientists to discover, try, and buy computing on demand, in the cloud. And he is the co-chairman of the International ISC Cloud & Big Data Conference series. Previously, he was Advisor to the EU projects EUDAT and DEISA, directed the 150 Million Euro German D-Grid Initiative, and was a member of the Board of Directors of the Open Grid Forum and of the US President's Council of Advisors for Science and Technology, PCAST.



Before, Wolfgang was a professor of computer science and mathematics at several universities in the US and Germany, and held leading positions at the North Carolina Grid and Data Center in Durham, Sun Microsystems in California, the DLR German Aerospace Center in Gottingen, and the Max-Planck-Institute for Plasmaphysics in Munich. In the 90s, he founded HPC software companies Genias and Gridware. The latter, which has been acquired by Sun Microsystems in 2000, developed the well-known distributed cluster workload and management system Grid Engine.

Abstract:

During the last two years UberCloud has developed Docker container technology into sophisticated HPC containers for engineering and scientific applications. Today, UberCloud Containers exist for applications like ANSYS Mechanical, Fluent, CFX, Workbench, Icepak, EM, LS-DYNA and LS-PrePost, for CD-adapco STAR-CCM+, for OpenFOAM, for Gromacs, and more in the pipeline. These application containers - 'develop once, run anywhere' - are portable on all in-house and cloud environments and allow seamless and transparent on-demand access to applications and data, with a look and feel identical to your workstation. These application containers are available on the online UberCloud Marketplace, with instant on demand access.

Hewlett-Packard

Title: HP's HPC Toolkit for Cloud Orchestration

Speaker: David Hanlon

Bio:



- BSc Engineering and Applied Sciences (Electronics)
- Nearly 33 Years experience in Industry, IT and complex manufacturing systems
- Production and Device Engineering (Advanced Technology – Semiconductor) National Semiconductor – 7 Years
- Advanced Manufacturing Systems and Automation – MES and HPC Motorola (18 Years).
- MES (3 Years) followed by HPC Specialization (4 Years - to date) – Hewlett Packard.
- Lead Architect HPCaaS Practice
- Married with 2 grown up kids and 2 not so grown up Golden Retrievers
- If he ever gets any spare time – photographs aircraft and builds model steam engines.

Abstract:

A method for integrating an existing or future traditional HPC/Grid environment with HP's Helion cloud infrastructure. This looks at some architectures that are being deployed by the HPCaaS team that allow clients to drive an existing or future client dedicated HPC/Grid installations from the Helion cloud stack and similarly integrate the HPC installation with the Helion cloud stack to allow resource sharing and management where that is appropriate to the client.

This presentation discusses the options being developed and deployed to support existing clients that take the form of dedicated (private) cloud, hybrid (dedicated) cloud and a full integration with a hybrid (shared) cloud, and the technical changes currently being delivered using the HPC Toolkit to enable this functionality.

Sardina Systems

Title: OpenStack Data Center Automation: Slashing Energy OpEx

Speaker: CJ Kenneth Tan

Bio: Dr Kenneth Tan is Director of Technical Solutions at Sardina Systems. He was previously with CloudFabriQ where he was the CTO, BNP Paribas where he served as Consultant in Group Risk Management, and OptimaNumerics where he was the Founder and CEO. He has previously developed and sold products and services to major customers in automotive, energy, defence, meteo, finance. He has also led commercial and technical teams in UK, Ukraine, Belarus, Estonia, Russia, Germany.

Abstract:

In many cloud data centers, servers utilization average between 6 to 12% -- the most efficient utilization barely touches 35%. When faced with a computing capacity problem, companies just put up more servers to cater to the workload, instead of running the machines the company already has at higher utilization. This short-term solution results in ever-increasing energy consumption and escalating costs of operations. Sardina Systems will be presenting FishDirector, an advanced solution for OpenStack environments, to lower energy OpEx, increase operations agility managing large scale facilities as one, increase RoI through increased server utilization, reduce servers and associated software CapEx.

ADVANIA

Title: HPC Cloud Projects in Iceland

Speaker: Per-Ola Svensson

Bio:

Do you know that Iceland is ranked as one of the world's best data center location? Iceland is especially suitable for the placement of datacenters for numerous reasons. One reason is that not a single second of downtime at Advania's data center is caused by natural disasters. And that is pure fact. A typical data center will use about as much energy for cooling as it takes to keep the servers and other equipment running. With a datacenter in Iceland, the customer takes advantage of so-called free natural cooling - this by using outdoor temperatures that never gets really hot or really cold. This is a key factor for achieving low average value of PUE. Thanks to Iceland's natural elements in combination with a specific datacenter design, Advania's data center is one of the most energy efficient datacenters in the world. Using the so-called "free-cooling" in Iceland, the power utilization is very efficient. From this datacenters Advania provides their HPC cloud services.

Abstract:

Altair Engineering

Title: Unlimited Innovation with Physical and Virtual Appliance Computing

Speaker: Ravi Kunju, Managing Director of Strategy and Marketing

Bio: Ravi Kunju brings over 20 years of experience in applying advanced numerical methods and analytics, specifically in HPC, to solve complex problems in the areas of computer-aided engineering (CAE) and business intelligence (BI). Ravi's career includes roles at Ford and Chrysler in the areas of crash-safety and advanced manufacturing (sheet metal forming), and at Altair in product design, software product management, and executive roles in global sales, regional management and strategic marketing. Ravi holds an M.S. in Mechanical Engineering from Wayne State University and an MBA from Ross School of Business, University of Michigan, Ann Arbor.



Abstract:

Cloud appliance computing democratizes the usage of cutting-edge CAE by reducing the barriers to HPC adoption for highly complex engineering tasks that typically require large-scale deployment and efficient usage of HPC clusters. Altair pioneered cloud appliances for HPC with their HyperWorks Unlimited private cloud appliance, offering companies a fully configured hardware and software solution for CAE including Altair's HyperWorks® suite and the award-winning PBS Works™ suite for HPC workload management. Now, with Altair's HyperWorks Unlimited Virtual Appliance, the offering goes fully into the cloud with an Amazon-hosted option that lets users get started with HPC in just minutes. This talk will provide an introduction to Altair's cloud appliance solutions which ensure users

have unrestricted access to the software they needed to explore and innovate -- making HPC far more accessible for the design community and paving the path to infinite exploration.

Tutorial E3/4: Insight CMU: Cluster Management Utility and Related Products

Hewlett-Packard

Title: Insight CMU: what you (may) know, what is new, what is next.
Featuring a live demonstration!

Speaker: Clément Poulain

Bio:

Abstract:

Altair Engineering

Title: Managing Complex HPC Systems with Insight CMU and PBS Professional 13.0

Speaker: Graham Russell, Technical Director Enterprise Computing EMEA, Altair Engineering

Bio: Graham Russell has 20+ years in technical and high performance computing working for various hardware and software companies. In his current role Graham has global responsibility for leading pre-sales and delivery engagements for complex projects and developing solutions requiring customised integrations with the PBS Works tools. Previously Graham was responsible for technical relationships with PBS Works partners and focused on large, complex and strategic PBS Works projects in EMEA. Prior to joining Altair Graham held various roles in pre-sales, product marketing, business development, and technical and region management.



Abstract:

In this presentation, Altair will provide an update of our CMU PBS Professional integration which automates the most commonly used cluster and workload management tasks, simplifying cluster deployment and management for users and administrators. We will also cover technical highlights from the new PBS Professional 13.0 release, so attendees can learn how Altair is moving our technology forward – the release includes with improved scalability, flexibility, usability and scheduling capabilities. The presentation will include a video demo of PBS-CMU integration features; it will also include tips and practical advice for managing a production cluster environment with Altair tools.

Adaptive Computing

Title: New Capabilities to Enhance HP CMU – Leveraging Nitro for HTC-specific Workloads

Speaker: Bernhard Schott

Bio:

Abstract:

HPC & Big Data SIG

Altair Engineering

Title: **Managing Big Data with Powerful Job Scheduling Plus Remote Visualization and Management Portals**

Speaker: **Bill Nitzberg**, PBS Works CTO

Bio: Dr. Bill Nitzberg is the CTO of PBS Works at Altair Engineering, Inc. With over 30 years in the computer industry, spanning commercial software development to high-performance computing research, Dr. Nitzberg is an internationally recognized expert in parallel and distributed computing.



Dr. Nitzberg has served on the board of the Open Grid Forum, he co-architected NASA's Information Power Grid, edited the MPI-2 I/O standard, and has published numerous papers on distributed shared memory, parallel I/O, clusters, job scheduling, and cloud computing. In his spare time, Bill tries to reduce his pack weight for his long-distance hiking trips.

This presentation will cover Altair's PBS Works suite for submitting, scheduling, monitoring, analyzing and visualizing the entire HPC workload lifecycle as it related to managing big data. With Altair's easy-to-use visualization portals, administrators have access to advanced job and license analytics to support data-driven planning and decision-making for their HPC infrastructures. PBS Works also powers public and private cloud implementations including Altair's own physical and virtual appliance offerings. In addition, Altair's recent release of PBS Professional 13.0 means even more performance and scalability to handle complex HPC requirements and big data needs. Attend this session to discover how this integrated suite, supported globally by the trusted leaders in HPC workload management, can drive efficiency, utilization and ROI increases in your organization.

Abstract:

AMD

Title: **AMD Visualization Features – An Update**

Speaker: **Karl Freund**

Bio: Karl Freund has been an executive in the IT hardware and software business for over 35 years, both at system vendors such as HP, Cray, and IBM, at Calxeda, the ARM Server Startup, and now at AMD as VP of High Performance Computing. Mr. Freund has been widely quoted in hundreds of articles, and has been a frequent speaker at technology and investment conferences. At AMD, he leads the company's strategy and execution of acceleration technologies that are enabling AMD to re-emerge as a leader in HPC.



Mr. Freund holds a bachelors degree from Texas A&M University in Applied Mathematics, and a Masters of Computer Science from the University of North Carolina at Chapel Hill.


Abstract:

Hewlett-Packard

Title: **Object Storage in HPC – an Overview of the Possibilities that the latest Storage Paradigms can bring in an HPC oriented Data Center**

Speaker: **Cedric Milesi**

Bio: Cedric Milesi has been working in High Performance Computing in HP for more than 10 years. He has participated in some of the biggest cluster deployment across Europe. He is now a Master Technologist in the HP EMEA central presale team focusing on hyperscale and parallel storage.



Abstract:

There is a new trend in the industry to replace traditional storage system with Object Store. Running on "off the shelf" servers, highly scalable, highly resilient those systems are being used in a range of case going from the massive storage system for online service to active archive. During this presentation we will detail some of the use such systems can have in a traditional HPC setup.

Cloud & Grid SIG

Adaptive Computing

Title: **Real world use case for Elastic Compute Shared Pools using Moab Elastic Compute/ Things Learned from Customers Running Grid and HPC Cloud-use Cases**

Speaker: **Paul Anderson, Bernhard Schott**

Bio: Growing and Shrinking HPC resources using a shared pool addresses many use-cases our customers are facing today. This presentation will give practical examples and lessons learned from one Medical Consortium's implementation of an Elastic Compute Shared Pool. From multiple independent, self-managed HPC environments to the economies of scale for the entire group, there are many reasons to implement an Elastic Compute Shared Pool.



Abstract:

Paul Anderson has been leading delivering custom HPC Solutions to address complex customer problems since 2007 and has led the integration teams for several Top500 clusters from build through acceptance. Since 2013, Paul has been leading the Professional Services team at Adaptive Computing to accelerate customer insights for HPC, Cloud and Big Data workflows.

Allinea

Title: **Improving Cluster Production with Allinea Tools**

Speaker: **Patrick Wohlschlegel**

Bio: Patrick graduated with an MSc from Bordeaux Graduate School of Engineering (ENSEIRB), specializing in parallelism and distributed systems. After being trained in Japan at the University of Kumamoto, he went to IBM in Paris as a pre-sales benchmarker. Through the experience he gained, he was then given the opportunity to create an HPC activity within a vendor reseller, Ovesys. As he developed the relationships between Ovesys and key HPC actors to design, build and deploy HPC centers, he naturally discovered Allinea. Today, Patrick is based in Warwick (UK) and works as a Technical Services Manager - sharing his expertise in HPC software tools to users all around the world.



Abstract:

HPC in the cloud brings lots of new questions and challenges. What nodes or types of nodes are most suitable for a particular application? Are data located in the right place? How is it possible to reduce runtime?

Over the last months, Allinea Performance Reports has been used successfully on a wide range of distributed and cloud facilities - including national multi-tiered datacenters and Amazon AWS. By enabling quick and efficient performance reviews, it has been invaluable to spectacularly reduce operating costs extremely quickly by improving application efficiency.

Nice

Title: **NICE DCV: Powerful Remote 3D Solutions enabling the HPC Cloud**

Speaker: **Karsten Gaier**

Bio: Dr. Karsten Gaier is working in the HPC space since 20 years offering sophisticated solutions for complex HPC, Grid and Cloud environments in CAE, O&G, Finance.



Abstract:

Remote 3D Visualisation enables the virtualisation of the workstation in the HPC Cloud and the creation of special purpose visualisation servers that sit alongside the computational HPC server.

A number of burning issues such as network overload, data security issues and high administration costs are causing many industrial and research organisations to strive towards this new model.

Such a deployment requires support for a mixed Linux and Windows user community possibly on a single server, the efficient sharing of hardware resources, and the flexibility for users to be able to run applications on physical or virtual machines running side by side on the same hardware.

NICE DCV is the only technology available today that ticks all these boxes and is helping customers move from the old model to the new HPC Cloud model faster than they thought possible.

Bright Computing

Title: **Manage Your HPC, Big Data & Private Cloud Infrastructure from a Single Pane of Glass**

Speaker: **Lee Carter**, Vice President EMEA

Bio:



Lee Carter joined Bright Computing in March 2014, bringing expertise in enterprise technical computing sales, to accelerate company growth across the EMEA region. Carter began his career at IBM, working with virtualisation and mainframe technologies, moving on to senior development, technical and sales roles, culminating as EMEA sales leader for grid and utility (cloud) computing. While at Microsoft and Adaptive Computing, Carter worked with clients both directly and through partners to leverage new technologies and enhance their HPC and cloud infrastructures. With degrees in Electrical Engineering and Computer Science from Purdue University, Carter's specialties include cloud computing, high performance Computing (HPC), virtualisation (mainframe & commodity), grid and utility computing

During the session, Lee Carter, VP EMEA at Bright Computing will expand on four case studies that showcase how Bright has positively impacted its global customers. For example, Carter will explore:

- Abstract:**
- How to expand your HPC cluster and be productive in minutes
 - How to deploy and manage multi-vendor HPC clusters from a central point of control
 - How to manage HPC and Hadoop from a single pane of glass
 - How to increase business agility and customer responsiveness with a Bright managed private cloud using OpenStack

CMU (Cluster Management Utility) SIG

Altair Engineering

Title: **Demo: Managing Complex HPC Systems with Insight CMU and PBS Professional 13.0**

Speaker: **Graham Russell**, Technical Director Enterprise Computing EMEA, Altair Engineering

Bio: Graham Russell has 20+ years in technical and high performance computing working for various hardware and software companies. In his current role Graham has global responsibility for leading pre-sales and delivery engagements for complex projects and developing solutions requiring customised integrations with the PBS Works tools. Previously Graham was responsible for technical relationships with PBS Works partners and focused on large, complex and strategic PBS Works projects in EMEA. Prior to joining Altair Graham held various roles in pre-sales, product marketing, business development, and technical and region management.


Abstract: In this presentation, Altair will present a video demo of our CMU PBS Professional integration which automates the most commonly used cluster and workload management tasks, simplifying cluster deployment and management for users and administrators. We will also cover technical highlights from the new PBS Professional 13.0 release, so attendees can learn how Altair is moving our technology forward – the release includes with improved scalability, flexibility, usability and scheduling capabilities. The presentation will include tips and practical advice for managing a production cluster environment with Altair tools.

Adaptive Computing

Title: **Best Practice Using Moab with CMU – Specific Customer Use Case**

Speaker: **Wil Wellington**

Bio: Wil Wellington is the Technical Director for Adaptive Computing in EMEA, specialising in Cloud, HPC and HTC.



Prior to joining Adaptive Computing Wil worked for HP Enterprise Services (EMEA) in the Defence and Security sector as a Infrastructure and Technical Consultant.

With a degree in Hardware and Software Engineering, Wil has worked for world leaders in financial and telecommunication companies focusing on next generation technologies and emerging markets.

Abstract:

HP Managed Services SIG

Hewlett-Packard

Title: Nucleus – HPC Managed Service

Speaker: Andrew Cusick

Bio:



- Degree – BSc Geology / Biology
- UK Programme manager – background in Government and Manufacturing Sectors at HP, PMP lecturer, Secretary of UK North Chapter of PMI
- Last 7 years – Programme Manager running services for HPC
 - Implemented major HPC Clusters internationally
 - This year – Consulting for : Astrium Satellites, Airbus, E.On, Tetra Pak, Deutsche Bank, ASML
- Lead for:
 - Global lead HPC Managed Service Offering
 - Business Development Consultant
 - Liaison with the UK Manufacturing Technology Centre UK

Nucleus – HPC Managed Services

Andrew Cusick – international consultant in HPC Services, and previously programme manager for HPC delivery at a major aerospace provider will outline the development of HP’s managed services offering. Drawing on real examples from recent successes in major blue chip industries in Industry and Finance, the presentation will focus on:

- The portfolio of Services HP is already delivering to its clients and how we do that
- How HP draws together a multi-disciplinary team to deliver HPC managed Services
- The benefits of HPC as a Service and how to overcome the pitfalls and obstacles to delivering excellent HPC services in the next generation of IT.
 - Delivering better total cost of ownership.
 - Improving job throughput.
 - Improving performance, reliability, security and delivery.
 - Freeing your intellectual capital to focus on the “day job”.
 - HP range of in house tools and capabilities.
- The future of genuine dial up / dial down services for professional IT – why it is so difficult for professional businesses to leverage HPC on a consumption based model and how this can be overcome.

Abstract:

If you are interested in managed services, or enhancing your own service to their existing users and in particular if you are intending to look at a consumption based model of HPC with variable compute capacity then this would be of interest.

Large System SIG Session (LS-SIG)

Title: Open Discussion of Progress, Suggestions, Issues and Problems

Speaker: