

OPENPOWER SUMMIT EUROPE AGENDA – OCTOBER 2016

THURSDAY 27TH OCTOBER – AGENDA OVERVIEW

	Restaurant Voramar	Mediterraneo + Foyer	Forum	Estrella de Mar	Princess 1
9:00 AM					
9:15 AM					
9:30 AM					
9:45 AM					
10:00 AM				Integrated Solutions Workgroup Chair: Randall Ross	System SW Workgroup Chair: Chris Austen
10:15 AM					Academia Discussion Group Chair: G. Narayanasamy
10:30 AM			Coffee Break		
10:45 AM					
11:00 AM					
11:15 AM					
11:30 AM					
11:45 AM					
12:00 PM					
12:15 PM					
12:30 PM	Lunch				
12:45 PM					
1:00 PM		Keynotes			
1:15 PM					
1:30 PM					
1:45 PM					
2:00 PM					
2:15 PM					
2:30 PM					
2:45 PM					
3:00 PM			Coffee Break		
3:15 PM					
3:30 PM		Keynotes			
3:45 PM					
4:00 PM					
4:15 PM					
4:30 PM					
4:45 PM					
5:00 PM					
5:15 PM					
5:30 PM					
5:45 PM					
6:00 PM					
6:15 PM		Reception Setup			
6:30 PM			Welcome Reception Demo Booths		Welcome Reception
6:45 PM					
7:00 PM		Welcome Reception Demo Booths			
7:15 PM					
7:30 PM					
7:45 PM					
8:00 PM					
8:15 PM					
8:30 PM					

FRIDAY 28TH OCTOBER – AGENDA OVERVIEW

		Mediterraneo + Foyer	Estrella de Mar	Princess 1	Forum	
7:00 AM						
7:15 AM						
7:30 AM	Coffee Break					
7:45 AM						
8:00 AM		System SW Chair: Jeff Scheel	Accelerator Workgroup Chair: Allan Cante	Performance Chair: Alex Mericas	Video Interviews	
8:15 AM						
8:30 AM						
8:45 AM		OpenPOWER for developers Chair: Randall Ross	Acceleration 1 Chair: Allan Cante	Physical Science Workgroup Chair: Andrea Bularelli		
9:00 AM						
9:15 AM						
9:30 AM	Coffee Break					
9:45 AM						
10:00 AM						
10:15 AM		Ecosystem and Solution Chair: Zaid Al-Ars	Acedemia Discussion Group Chair: D. Pleiter	OpenSTACK and OpenPOWER enablement and demo Chair: Marcelo Perazolo		
10:30 AM						
10:45 AM						
11:00 AM						
11:15 AM						
11:30 AM						
11:45 AM	Coffee + Cookies					
12:00 PM		Panel Business in an Open world Moderator: Calista Redmond		Acceleration 2 Chair: Bruce Wile		
12:15 PM						
12:30 PM						
12:45 PM						
1:00 PM	Summit Ends					

13:00 OPENING INTRODUCTION

Calista Redmond *OpenPOWER Foundation President*
John Zannos *OpenPOWER Foundation Chair*

13:15 OPENPOWER ROADMAP

Rani Borkar *IBM Vice President of OpenPOWER Development*
Steve Fields *IBM Fellow and Chief Engineer Power Systems*

13:30 THE IBM - BARCELONA SUPERCOMPUTING DEEP LEARNING CENTRE

Mateo Valero *Director, Barcelona Supercomputing Centre*

Announcing the creation of the joint IBM-BSC Deep Learning Centre. This centre will provide the framework for conducting joint research and development projects on the Deep Learning domain, an essential component of cognitive computing, with focus on the development of new algorithms to improve and expand the cognitive capabilities of deep learning systems. Additionally, the centre will also research on flexible computing architectures –fundamental for big data workloads– like data centric systems and applications.

13:45 SMARTER INNOVATION AT SCALE AT THE HARTREE CENTRE

Neil Morgan *Future Technologies, STFC Hartree Centre*

The STFC Hartree Centre, as a collaboration with IBM Research, is focused on research into the application of cognitive and data-centric technologies to transform the capabilities of UK industry. The talk will give a brief overview of the centre from a technology perspective and give examples of high impact uses of IBM Watson and OpenPOWER technology.

14:00 ATOS - IBM PARTNERSHIP FOR POWER

Gilles Delamplé *Director, Global Escala*

When partnership rhymes with success and long term

14:15 BLUEBEE: HIGH PERFORMANCE GENOMICS-AS-A-SERVICE USING ELASTIC ACCELERATION

Zaid Al-Ars *Co-founder Bluebee*

This presentation describes the NGS analysis platform of Bluebee that enables scalable, accelerated and secure private-cloud solution targeted at clinical application of NGS analysis techniques. Scalability of the platform allows for seamless increase of the throughput when required by the clinic, which enables more patients to be analysed per hour. Acceleration, on the other hand, allows for reducing the latency of the analysis, which enables faster sample to diagnosis for a specific patient. Security of the platform makes it viable for clinical users, where privacy is assured by regulation. The presentation also discusses the results of a couple of case studies carried out in practice using the Bluebee platform for executing cancer diagnostic pipelines, showing a reduction of execution time from 18 days to only 2 days by scaling up on a large cloud infrastructure.

14:30 GPU-ACCELERATED DATABASE: CONVERGENCE OF BUSINESS ANALYTICS & AI

Charles Sutton *Managing Director, Europe, Kinetica*

In this session, we will introduce you to one of the latest technology trends in real-time data processing, deep learning, and data discovery. Charles Sutton will demonstrate the power & real-time processing capabilities of the GPU-accelerated database, exploring how this technology can be used to solve real-world business & analytics problems for all types of data in motion. Some of the use-cases discussions will include how this technology integrates with popular open source solutions like Kafka, Apache NiFi, H2O, Hadoop to create Hyper-converged real-time data infrastructure; he will also provide comparison to traditional EDW, MPP type solution and in-memory databases like SAP Hana or Oracle TimesTen.

14:45 OPENSTACK OUT IN THE OPEN

Ildiko Vancsa *Ecosystem Technical Lead, OpenStack Foundation*

OpenStack is not just an open source cloud platform but also a large community at the same time. The presentation will describe the main concepts behind the software package with case studies while also highlighting the importance of open collaboration.

15:00 DEPLOYING AND OPERATING OPENSTACK CLOUDS ON OPENPOWER

Edward Shvartsman *Senior Technical Staff Member, Power Systems, IBM*

15:15 SWIFT OBJECT STORE DEPLOYMENT ON OPENPOWER

Jacob Caspi, *Principal Systems Architect, AT&T*

Christian Reis *VP Hyperscale, Canonical*

Tom Mathews *Distinguished Engineer, Power Systems, IBM*

AT&T will review a use case for geo-distributed Swift object store deployments at scale. With its Domain 2.0 effort, AT&T is a big proponent of open source software and hardware. You will hear about AT&T's Proof of concept testing with IBM and OpenPOWER to provide an open hardware and open source software object storage solution with benefits in density and performance.

15:30 BREAK

16:00 ANNOUNCEMENT

Prof. Krcmar *TU Munich*

16:15 THE POWER OF OPEN -- OPEN PLATFORMS SHAPE THE FUTURE

Georg Greve *CEO, Kolab & Founder of Free Software Foundation Europe*

Information Technology is currently in a crisis of confidence. For decades, security professionals have preached radical openness as the only path toward establishing trust and confidence. Now, OpenPOWER is in a position to deliver just that. Openness from the bottom that reaches as high as the user demands. And because OpenPOWER also delivers value and technical innovation beyond the state of the art, it is likely to be at the core at the next major evolution of the IT industry and economy.

16:30 INTRODUCING THE OPENPOWER AMBASSADOR PROGRAM

Randall Ross *Ubuntu Community Manager, Canonical*

Are you passionate about OpenPOWER? This session is a briefing on the Ambassador Program: What it is, How to participate, and why you should!

16:45 DEVELOPER FOCUSED ANNOUNCEMENTS

Bruce Wile *CAPI Chief Engineer, IBM*

John Zannos *OpenPOWER Foundation Chair*

We are announcing an open source, OpenPOWER project that provides an acceleration platform for programmers, called "CAPI SNAP Framework" (Storage, Networking, and Analytics Programming". The framework makes it easy for developers to create specialized accelerated algorithms in C++, Go and other high level programming languages-- utilizing the leading-edge CAPI technology.

Plus announcement of the winners of the OpenPOWER Developer Challenge

17:00 FAST FORWARD EVOLUTION

Fabrizio Magugliani *E4 Computer Engineering*

Standard C.O.T.S. are known to thrive and proliferate in HPC environments, where price, performance and power consumption are the crucial issues that cannot be overlooked. However as technology's leap forward is getting faster than ever, users are demanding extra performances, higher productivity and more variety of solutions to choose from.

With its usual technology agnostic approach that is at the core of our success, E4 was one of the first to introduce an OpenPOWER technology based platform aimed at the most demanding HPC workloads.

During the presentation we will illustrate the reasons behind our choice to adopt this technology with a number of examples in terms of performance, ease of use and applications.

17:15 EARLY INSIGHTS WITH NVLINK SYSTEMS

Danny Kiernan *IBM Senior Alliances Manager, NVidia*

17:30 ACCELERATING DEEP LEARNING TRAINING WITH POWER8, P100, AND NVLINK

Scott Soutter *Product Manager, Deep Learning and Accelerated Databases, IBM*

Overview of the upcoming IBM software distribution for deep learning, including optimization for NVIDIA P100 and NVLink.

17:45 NEXT GENERATION PERFORMANCE AND SCALABILITY WITH MELLANOX "SMART" INTERCONNECT AND OPENPOWER

Chloe Ma *Senior Director, Cloud Marketing, Mellanox*

Pushing the frontiers of science and technology will require extreme-scale computing with machines that are 500-to-1,000 times more capable than today's supercomputers. Technology development has had to keep up in order to enable such performance leaps. Today and for the future this includes extreme innovation with the OpenPOWER architecture. We will cover the newest technology advancements in networking and is part of the industry's co-design efforts including how the network is becoming a co-processor of the POWER-based system with Mellanox SHArP technology, 100Gb/s networking with CAPI today and future capabilities with 200Gb/s for 2017.

18:00 OPENPOWER WORKGROUP UPDATE

Jeff Brown *Technical Steering Committee Chair*

WORKING SESSIONS

PERSONALISED MEDICINE WORKGROUP WED 26TH OCT 15:30 – 17:00 ESTRELLA DE MAR

Lead: Zaid Al-Ars (Bluebee)

This WG is a persistent group that will have the application domain of personalized medicine as its focal point. It will analyze commonly used, computationally intensive personalized medicine applications, identifying specific bottlenecks faced by these applications. Ensuring efficient execution of these applications will be addressed in two ways. First, optimized solution architecture configurations and requirements will be proposed. Secondly, optimizations to the code of these applications will also be investigated and outlined.

INTEGRATED SOLUTIONS WORKGROUP THURS 27TH OCT 09:00 – 10:30 ESTRELLA DE MAR

Lead: Randall Ross (Canonical)

This WG is a persistent group that will provide identification of Solution Domains and Solution Sets within those domains and provide a detailed description of the solution set and its potential opportunity relative to the OpenPOWER Foundation and the OpenPOWER ecosystem. Solution Sets are targeted and solve/address specific problems within a Solution Domain. They may consist of combinations of applications, accelerators, enhanced networking, etc. The WG will also evangelize creation of solution set targeted workgroups and monitor the state of the solution sets relative to ongoing evolution and to evangelize and advocate for re-energizing if appropriate

SYSTEMS SOFTWARE WORKGROUP THURS 27TH OCT 09:00 – 09:45 PRINCESS 1

Lead: Chris Austen (IBM)

The Systems Software work group will serve as the stewards, incubators, and drivers (when feasible) of system software enablement for OpenPOWER Foundation hardware. As such, the work group will work to ensure the availability of all software required to boot, run, and manage Linux on OpenPOWER compliant systems. This includes platform firmware, virtualization environments, Linux itself, toolchain (compilers, core libraries, debuggers, etc.) and platform management software. This does not include software specifically used to enable end user application execution. The goal for all software is availability under an appropriate open source license.

Presentations related to this workgroup are held on Friday 28th Oct, 08:00 – 08:45 @ *Mediterraneo*

ACADEMIC DISCUSSION GROUP THURS 27TH OCT 09:45 – 10:30 PRINCESS 1

Lead: Ganesan Narayanasamy (IBM)

The Academia Discussion Group (ADG) within the OpenPOWER Foundation is an academic forum, which mainly targets academic members of the foundation but is open for all other members. The goals of ADG are to provide training and exchange of experience and know-how, to provide a platform for networking among academic members, to work on the engagement of the HPC community and to enable co-design activities.

Presentations related to this workgroup are held on Friday 28th Oct , 10:15 – 11:45 @ *Estrella de Mar*

TECHNICAL STEERING COMMITTEE MEETING THURS 27TH OCT 11:00 – 12:30 PRINCESS 1

Lead: Jeff Brown (OpenPOWER Foundation Technical Steering Committee Chair)

Please note that this meeting is by invitation only

ACCELERATOR WORKGROUP FRI 28TH OCT 08:00 – 08:45 ESTRELLA DE MAR

Lead: Allan Cantle (Nallatech)

The Accelerator work group will be a persistent, standing group that defines, documents, manages, and maintains standards which define the interfaces between the processor and accelerator devices and their associated development tools. This includes hardware, firmware and any other software require for accelerators to operate within OpenPOWER compliant systems. In this session the discussion on Accelerator Projects and the workgroups forward direction will be outlined.

Presentations related to this workgroup are held on Friday 28thOct, 8:45 – 09:30 @ *Estrella de Mar* and 12:00-13:00 @ *Princess 1*

PHYSICAL SCIENCE WORKGROUP

FRI 28TH OCT

08:45 – 09:30 *PRINCESS 1*

Lead: Andrea Bulgarelli (INAF, L'Istituto Nazionale di Astrofisica)

This Work Group aims at addressing the challenges of Physical Science projects by developing use cases, identifying requirements and extracting workflows, to better understand common workflows with related needs and pain points. Based on these identified use cases, the members of the Work Group discuss the definition of reference solutions which provide broad value to the physical sciences users and that will be made public. Working around use cases, the WG allows the OpenPOWER Foundation to be a forum between scientists and technical solutions developers, and also between scientists of different fields and projects to share experience and solutions with each other.

PANEL SESSION

"BUSINESS IN AN OPEN WORLD"

FRI 28TH OCT

12:00 – 13:00 *MEDITERRANEO*

Moderator: Calista Redmond, President OpenPOWER Foundation

How can a business exist and thrive in an Open-* world? Come and hear from our panel of business leaders with a huge combined experience in this area and what it takes to succeed.

Panel members:

- Randall Ross – Canonical
- Jeff Scheel – IBM
- Georg Greve – Kolab
- Major Hayden – Rackspace
- Rob Taylor – Reconfigure

PRESENTATIONS

SYSTEMS SOFTWARE

FRIDAY 28TH OCT

08:00 – 08:45 *MEDITERRANEO*

Session chair: Jeff Scheel (IBM)

OpenBMC, A Reference Firmware Stack

Chris Austen (IBM)

OpenPOWER is open with the exception of the BMC. That all changes with the introduction of OpenBMC. This talk will focus on what it is and how the community organized team of developers are using it on many OpenPOWER servers along with the future directions including P9, OpenStack and OpenCompute. Audience: OpenStack Ansible and Ironic developers, Server Design architects, System Management architects, Manufactures, Cloud providers. People will leave the talk wanting to get engaged with the OpenPOWER Foundation workgroups knowing they have a voice

Measuring and Managing Power Consumption

Todd Rosedahl (IBM)

This presentation will include an overview of the power, thermal, and performance data that can be collected from OpenPOWER servers via various methods, including a newly open sourced profiling tool called AMESTER. The power/performance knobs, such as processor frequency, that are under the control of the On Chip Controller (OCC) will be described and the overall OCC power management functions will be highlighted.

Ensuring POWER9 Success with System-level Simulation for Early Software Verification

Saif Abrar (IBM)

Success of POWER9 systems depends on timely availability of various software stacks. Absence of suitable hardware poses a verification challenge for new firmware and drivers. This presentation describes a behavioral system-level simulation environment, made of a POWER simulator and behavioral models of devices and interconnects; these are used for early verification of the software. Various product-level software like Linux, OPAL, IBM Power Hypervisor drivers, IBM Flexible Service Processor firmware, SBE stack, Hostboot driver and P7/P8 stacks have been debugged and verified successfully before actual HW, shortening the HW bringup cycle to save cost and time. It is versatile to also validate the post silicon tests for HW IPs in the system.

PERFORMANCE

FRIDAY 28TH OCT

08:00 – 08:45 *PRINCESS 1*

Session chair: Alex Mericas (IBM)

OpenPOWER Performance

Alex Mericas (IBM)

OpenPOWER builds on the strengths of the Power Architecture through open collaboration. We will discuss how IBM enables partners to exploit the performance of OpenPOWER offerings. This session will introduce the performance characteristics of the 1-Socket and 2-Socket reference designs and give an overview of performance results for these systems. To support OpenPOWER Partners IBM provides Performance Enablement, which will be discussed.

IBM XL C/C and Fortran Compilers for OpenPOWER

Shereen Ghobrial (IBM)

IBM XL C/C and Fortran compilers are built on an industry wide reputation for robustness, versatility, standards compliance and performance. This presentation will provide the latest update on IBM XL compilers for OpenPOWER, which will cover major features to enhance portability, compatibility and performance. XL C/C for OpenPOWER is built with Clang front end components and IBM highly optimizing backends. It provides improved GCC compatibility and language standards support for easier migration and enhanced capability. This presentation will introduce IBM XL C/C and Fortran support for GPU exploitation through CUDA and OpenMP accelerator extension, and address how applications can benefit from the IBM XL compilers' optimization.

Emerging Workload Performance Evaluation on Future Generation OpenPOWER Processors

Saritha Vinod (IBM)

Open source innovation has resulted in rapid emergence of many significant real-world workloads. Next generation processor design need to evaluate the emerging workload behavior to ensure optimal performance. Hence emerging workload traces capturing the essence of these applications are essential to gather insights on key characteristics and allow performance evaluation in future generation of processors. There are unique challenges in generation of traces and using them for evaluation. How do we ensure that the collected trace is truly a representation of the workload behavior? What kind of performance evaluations can be done using these traces? In this talk we present about various techniques and challenges in generating representative workload traces and how they enable performance evaluation for future generation of processors.

OPENPOWER FOR DEVELOPERS

FRIDAY 28TH OCT

08:45 – 09:30 *MEDITERRANEO*

Session chair: Randall Ross (Canonical)

OpenPOWER: A Community Tour

Randall Ross (Canonical)

An overview of the OpenPOWER community for ISV's (and developers). If you're thinking about OpenPOWER and need resources and inspiration, this talk will highlight what's available, what's happening, and why OpenPOWER is the "next big thing".

The CAPI SNAP Framework for Programmers

Bruce Wile (IBM)

Since the release of the CAPI Technology and associated Developer Kit in 2014, the realm of FPGA acceleration has been dominated by those with the computer engineering skills of VHDL and Verilog. We are announcing an open source, OpenPOWER project that provides an acceleration platform for programmers, called "CAPI SNAP Framework" (Storage, Networking, and Analytics Programming). The framework makes it easy for developers to create specialized accelerated algorithms in C++, Go and other high level programming languages-- utilizing the leading-edge CAPI technology. This session will briefly describe the framework, the "acceleration action" programming paradigm, and the simple APIs used to invoke accelerated actions.

How to become a SuperVessel / Developer Cloud?

Yong Hua Lin and Peter Hofstee (IBM)

This brief presentation will cover Supervessel, an OpenStack-based OpenPOWER cloud accelerated with GPUs and FPGAs. Infrastructure available for academia and developers in China, and the US, and also a recent addition in Europe will be highlighted. A number of examples of Supervessel-based projects that leverage GPU and FPGA acceleration will be covered. Finally we explain how to get started if you not only want to use Supervessel, but plan to build one of your own!

ACCELERATION 1

FRIDAY 28TH OCT

08:45 – 09:30 *ESTRELLA DE MAR*

Session chair: Allan Cattle (Nallatech)

A 101 Guide to Heterogeneous, Accelerated, Data Centric Computing Architectures

Allan Cattle (Nallatech)

Data Centric terms have recently become popularized in the world of computing, but this shift in lexicon has not really been well communicated. This presentation aims to provide a 101 guide to the disadvantages of traditional Von Neumann, CPU Centric Architectures and their 25 year history of "patching up" and "working around" the well understood Von Neumann Memory Wall Syndrome with increasingly complex caching hierarchies and structures. The presentation will close with highlighting how IBM OpenPOWER members are all feverishly collaborating to make this seismic shift in the computing industry happen as smoothly and quickly as possible to bring far more computational horse power for far less energy & cost become a true reality.

ConTutto - A flexible memory interface in the OpenPOWER ecosystem

Thomas Roewer (IBM)

ConTutto is a configurable platform for innovation in the memory subsystem of an IBM Power System node. By replacing the memory buffer ASIC of a commercial IBM Power server with a high-end FPGA connected to the processor's memory link, we enable multiple system-level innovations including 1) Systems Research with new memory technologies, such as MRAM and

NVDIMMs, and 2) Near-memory acceleration of data flowing between processor and memory. In this talk, we will present details of the ConTutto platform, and show system-level implementations of both use-cases. This work is partially funded through the Department of Energy FastForward-2 contract.

Everyone agrees: computers will stop getting faster or cheaper. What do we do now?

Rob Taylor (Reconfigure)

In a world increasingly driven by data and analysis, the future depends on the cost and efficiency of computing. Technology entrepreneur Rob Taylor shows the limits of current designs and a possible path forward using hardware acceleration.

ACCELERATION 2

FRIDAY 28TH OCT

12:00 – 13:00 PRINCESS 1

Session chair: Bruce Wile (IBM)

The CAPI SNAP Framework Deep Dive

Bruce Wile (IBM)

Since the CAPI Developer Kit debuted in November, 2014, we've seen multiple infrastructure and ecosystem enablement contributions. These, in turn, have fostered the growth in the number of CAPI accelerators. This presentation will take a quick look back at the ecosystem growth before focusing on the next generation CAPI SNAP^{*)} Framework being introduced at the OpenPOWER Summit Europe. The framework makes it easy for developers to create specialized accelerated algorithms in C++, Go and other high level programming languages.

^{*)} Storage, Networking, and Analytics Programming

Leveraging OpenPOWER and FPGAs to accelerate machine vision and learning

Lisa Liu (Xilinx)

Leveraging coherent memory integration via IBM's CAPI interface, IBM POWER8 CPUs combined with Xilinx FPGA accelerators provide higher compute efficiency, lower power utilization and total cost of ownership for critical workloads in data centers and HPC. In this session, we will provide an overview of how FPGA acceleration can enhance POWER systems for a broad range of applications with a particular emphasis on machine vision and learning workloads such as image recognition and video scaling. We include a theoretical study, leveraging UC Berkeley ParLab Roofline models, as well as present some recent results from a selection of prototypes on a POWER8 system.

New Memory Controller and Very Low Latency Network for Accessing Low Latency (200-400ns) NVM Storage

Luiz M Franca-Neto (HGST)

Current Datacenters use Ethernet/IP fabric and switches to access NAND-Flash based local and remote storage. 75us readout latency NAND-Flash are accessed through additional 5us or 10us Ethernet switch latencies in the datacenter fabric. However, with emerging low latency NVM-based storage with readout latencies ~1,000x times shorter, at ~200ns, Ethernet/IP fabric and switches are clearly inadequate. We present and discuss a new very low latency fabric (~100ns) topology of critical local and remote access (RMA) for these emerging NVM storage solutions. We outline the necessary and viable changes in the design of OpenPOWER memory controller to enable an RMA service launched from WRITE/READ commands on the memory bus.

Phase Change Memory Access in OpenPOWER Systems using CAPI

Nikolaos Papandreou (IBM)

Novel forms of storage class memory (SCM), such as phase change memory (PCM), promise low latency and small granularity of R/W access while offering high endurance and storage density. In this work we demonstrate low latency PCM access in OpenPOWER systems via CAPI. Our platform consists of OpenPOWER servers (IBM S824L, Tyan Palmetto) equipped with CAPI-enabled FPGA cards and custom non-volatile DIMMs made of PCM chips. In the talk, we will give an introduction to PCM as the top contender for SCM and present the architecture of our prototype system and FPGA-based CAPI-enabled PCM controller. Moreover, we will present latency and performance measurements using our custom PCM DIMMs as well as custom HW that emulates future PCM technology.

Session chair: *Randall Ross (Canonical)*

OpenPOWER Ready

Jeff Brown (IBM)

OpenPOWER Ready™ is a mark used by the OpenPOWER Foundation to enable ecosystem product developers to indicate their product has been shown/demonstrated to meet a minimum set of characteristics and should be interoperable with other OpenPOWER Ready products. Come and hear more about this program and how your organisation can be involved.

Increasing the Density of Apache Spark in the Cloud Through Coherently Attached Flash on POWER8

Thomas Hubregtsen (IBM)

Traditionally, disk bandwidth has been the limiting factor in Big Data processing systems. In Spark, this problem is alleviated by using in-memory computation, thereby providing the user with a high amount of bandwidth. Unfortunately, memory has become the most expensive part of a computer system which typically limits the available capacity. Spark starts spilling to disk once the available memory is exhausted which can have a significant impact on the performance. In IBM Research we aim to leverage user-addressable Flash storage as a cheaper alternative to increasing the memory while providing better performance compared to hard disk drives. We implemented an initial version of Spark that uses Flash attached through the CAPI in POWER8.

Highly Scalable OpenPOWER Big Data Framework for DNA Analysis Pipelines

Zaid Al-Ars (Bluebee)

OpenPOWER vendors and users as well as Operations/DevOps tool providers in our community can benefit and further engage to provide feedback and collaborate in this effort.

Deep Learning on Power

Samuel Cozannet (Canonical)

Artificial intelligence will be a core component of a large majority of future applications. We present a reference architecture to start training neuronets in minutes on Power, as well as an example workflow on an example Network Intrusion Detection application, from gathering the training data, to performing training and finally moving the pre trained model to production.

Low Latency Edge Supercomputers for Connected Vehicles and Video Analytics Using OpenPower and RapidIO

Devashish Paul (IDT)

The main reason for Mobile Edge Computing is it place computing resources close to end users to reduce the round trip latency between transaction initiation and application success without congesting 4G and 5G networks. Latency is not only limited to physical distance, but also the logical “distance” between computing elements inside the Edge Computer, which can be the difference between application success and failure, or as minimum, unsatisfactory delivery of the MEC enabled use case. As both Operators and Hyperscale Cloud Data Center owners try to work together to deliver more optimized network configurations to enable these use cases, in this presentation we will discuss low latency solutions at the edge. The presentation covers HPC like low latency heterogeneous computing appliances that can be placed at the wireless network access point, in the C-RAN or in the Central office, enabling a variety of applications that will be dependent on low latency computing using IBM Power CPUs, accelerators and RapidIO interconnect. The presentation covers the Connected Vehicle offload project that IDT is conducting at 5G Lab Europe where the MEC solution is built to manage vehicles and fleet data in a cell coverage area. The presentation also covers the real time object recognition and video analytics use case as well as broadcast video with Edge Computing

POWER9 chip technology

Jeff Stuecheli (IBM)

This presentation describes features of the POWER9 family, with optimizations for one- and two-socket servers, DDR4 direct-attached memory, PCI Express Gen 4, and multi-socket enterprise servers with robust buffered memory systems. The design includes a new core microarchitecture featuring execution-slice technology, delivering enhanced performance with improved execution efficiency. The new design includes intelligent features for optimizing cloud workloads and improves virtualization efficiency. It features a robust set of accelerated heterogeneous computing technologies, enabling innovative solutions in the OpenPOWER ecosystem.

Lead: Prof Dirk Pleiter (Jülich Supercomputing Centre)

The Academia Discussion Group (ADG) within the OpenPOWER Foundation is an academic forum, which mainly targets academic members of the foundation but is open for all other members. The goals of ADG are to provide training and exchange of experience and know-how, to provide a platform for networking among academic members, to work on the engagement of the HPC community and to enable co-design activities. In this session we will give an overview on some of the ADG activities, provide information on the IBM academia program, present selected high-lights from using OpenPOWER technologies and architectures for science and close with an outlook on future opportunities from a supercomputing centre's perspective.

Agenda:

- | | |
|---|--|
| 1. Introduction | Dirk Pleiter (JSC, Germany) |
| 2. IBM academic program | Presenter: Ganesan Narayanasamy (IBM, India) |
| 3. Applications from physical sciences | Presenter: Andrea Bulgarelli (INAF, Italy) |
| 4. GPU-accelerated POWER for supercomputing | Presenter: Dirk Pleiter (JSC, Germany) |
| 5. Future opportunities using OpenPOWER | Presenter: Sadaf Alam (CSCS, Switherland) |

Session chair: Marcelo Perazolo (IBM)

In this 90-min combined session we will showcase 5 different ready-for-production scenarios where we exploit the advantages of OpenPOWER with OpenStack and related technologies such as Swift, Ceph, Trove DBaaS, Private Clouds and Operational Management aided by Open DevOps solutions, such as Availability & Health Monitoring with Nagios and Log Collection & Analysis with the ELK Stack. We will show how any user can deploy these applications themselves with readily available Open Source resources that focus on providing full support for OpenPOWER Hardware Reference Architectures. Different architectures will be discussed and explained how to best taken advantage of the Power architecture superior data processing capabilities. In addition, we will contrast two different approaches for the deployment of OpenStack clusters - an open approach with community support, based on OpenStack-Ansible, and a commercial approach, based on RedHat's RDO distribution using PackStack.

All scenarios demonstrated will discuss architecture decisions and the advantages acquired by the use of OpenPOWER hardware in diverse roles, such as Compute and Storage nodes. Most of these scenarios allow for flexible deployment with Ansible playbooks, providing easy configuration and maintenance with full support for all necessary packages to run on the Power architecture on top of leading Power Linux OSs provided by Canonical and RedHat.

Agenda:

- | | |
|--|--|
| 1. Deploying Private Clouds using Open Software on OpenPOWER | Presenter: Kyle Henderson, IBM Systems |
| 2. OpenPOWER Ceph and Swift Storage Solutions | Presenter: Sam Matzek, IBM Systems |
| 3. Managing OpenPOWER & OpenStack Operations with Nagios and ELK | Presenter: Marcelo Perazolo, IBM Systems |
| 4. An Open DBaaS Architecture and Implementation for OpenPOWER | Presenter: Nikhil Hegde, IBM Systems |
| 5. Build an Open Cloud with RDO on PowerLinux | Presenter: Xiandong Meng, IBM Systems |