Communiqués de presse

Annonces majeures d’IBM dans le cadre de la conférence OpenPOWER Summit

Paris - 07 avr. 2016: Les centres informatiques connaissent actuellement un basculement majeur : le déluge de données pousse les clients à rechercher des solutions informatiques alternatives optimisées pour répondre à la demande croissante d'analyse de mégadonnées. Face à ce besoin, les leaders de l'industrie informatique ont compris que les systèmes informatiques généralistes et propriétaires ne sont pas la réponse attendue. Le développement collaboratif est la solution.

Cette semaine, lors de la conférence OpenPOWER Summit, les membres de l'OpenPOWER Foundation bouleversent à nouveau le status quo en introduisant de nouvelles technologies et de nouvelles implémentations dans les centres informatiques, fruits de la collaboration ouverte entre ses membres.

Les membres de l'OpenPOWER Foundation unissent leurs forces à celle de l'Open Compute Project pour annoncer de nouvelles innovations. En effet, ces deux groupes partagent la même conviction : l'ouverture permet de multiplier les choix et améliorer la technologie. L'Open Compute Project vise à partager des concepts et des techniques de conception de centres informatiques, tandis que l'OpenPOWER Foundation se focalise sur le développement ouvert de serveurs et de systèmes informatiques.

Les annonces majeures de cette semaine sont les suivantes :

- Google et Rackspace annoncent le développement en commun d'une architecture de serveur basée sur le processeur IBM POWER9.
- Rackspace annonce ses plans de déploiement de son serveur « Barreleye », le premier serveur conjuguant OpenPOWER, Open Compute et Open Stack.
- IBM ajoute à sa gamme Linux on POWER des serveurs conformes à l'Open Compute Project, pour supporter des applications big data / analytiques ainsi que cognitives en mode Cloud.
- L'Université du Michigan annonce un projet de déploiement d'un système OpenPOWER pour le calcul intensif et la recherche scientifique.

Ces annonces sont complétées par plus de 50 autres innovations couvrant tout le spectre des besoins des centres informatiques. Les membres de l'OpenPOWER Foundation annoncent également les solutions suivantes :

- DRAGEN Genomics Platform, une nouvelle appliance développée par Edico Genome, IBM et Xilinx, permettant une analyse ultra-rapide de données génomiques.

- Plus de 12 nouveaux accélérateurs matériels, développés entre autre par BittWare, IBM, Mellanox et Xilinx, augmentant ainsi l'usage de la technologie d'accélération CAPI.
- TYAN GT75-BP012, un serveur POWER8 de taille 1U, présentant d'exceptionnelles capacités pour le calcul en mémoire vive (in memory computing).
- Une deuxième génération de serveurs IBM OpenPOWER HPC (High Performance Computing), développés en collaboration avec NVIDIA et Wistron, équipés de processeurs POWER8 connectés directement à des GPU NVIDIA Tesla à travers l'interface haute performance NVLink.

###

OpenPOWER Ecosystem Propels Open Innovation in Hyperscale Data Centers
Google and Rackspace Develop OpenPOWER System for The Open Compute Project;
IBM Announces Intent to Expand Line of POWER-based Scale-out Linux Servers

OPENPOWER SUMMIT, San Jose, Calif. – April 6, 2016 – The OpenPOWER Foundation, a consortium of more than 200 leading technology companies, organizations and individuals innovating around the POWER processor, today announced more than 50 new open innovations to help companies better solve grand challenges around big data.

Many new community innovations unveiled today are designed to be incorporated into the Open Compute Project product portfolio.

Among these, Google, a founding member of the OpenPOWER Foundation, announced today that it is developing a next-generation OpenPOWER and Open Compute Project form factor server. Google is working with Rackspace to co-develop an open server specification based on the new POWER9 architecture, and the two companies will submit a candidate server design to the Open Compute Project.

Additionally, Rackspace has announced that “Barreleye” has moved from the lab to the data center. Rackspace anticipates “Barreleye” will move into broader availability throughout the rest of the year, with the first applications on the Rackspace Public Cloud powered by OpenStack. Rackspace and IBM collectively contributed the “Barreleye” specifications to the Open Compute Project in January 2016. The specifications were formally accepted by the Open Compute Project in February 2016.

“*Today’s IT infrastructure leaders seek open technology alternatives to processor and system architectures,”* **said John Zannos, Chairman of the OpenPOWER Foundation, and Vice President of Cloud Channels and Alliances, Canonical**. “*Customized solutions and open building blocks are quickly becoming required options for system design. Collaborative innovation, the hallmark of both the OpenPOWER*

Foundation and the Open Compute Project, is essential to building the next generation data center."

"*We're thrilled to take the next step in our work with the OpenPOWER and Open Compute Project communities,*" **said Maire Mahony, Hardware Engineering Manager, Google, and OpenPOWER Foundation Board Member**. "*We are committed to open innovation, and to optimizing performance and cost in data centers. Working with Rackspace, we will submit a POWER9 server design to the Open Compute Project that will address the diverse requirements of end customers for data center services.*"

"*We are excited to work with Google on our POWER9 OpenPOWER-based, Open Compute Project form factor server*," **said Aaron Sullivan, Open Compute Project Incubation Committee Member and Distinguished Engineer at Rackspace**. "*OpenPOWER processors combined with acceleration technology are fundamentally changing server and data center design today and into the future. OpenPOWER provides a great platform for the speed and flexibility needs of hyperscale operators as they demand ever-increasing levels of scalability.*"

"*Our ongoing work with the OpenPOWER Foundation is a natural extension of our commitment to open collaboration and innovation in data center technology,*" **said Amber Graner, Director of Operations, Community Manager, the Open Compute Project**. "*The Open Compute Project is focused on efficiency, flexibility, and openness—and we recognize the importance of the POWER processor and the robust OpenPOWER ecosystem for the future of sever design.*"

The Open Compute Project is a member of the OpenPOWER Foundation Advisory Group.

IBM Expands its Linux-only Portfolio Leveraging OpenPOWER Innovation

IBM announced that it plans to add systems to its LC line of servers. The LC line, launched in October of 2015, infuses OpenPOWER technology into IBM's scale-out server lineup. As a result of dozens of proof of concepts in areas like hyperscale data centers, high performance computing and large enterprises, IBM intends to make the following additions to the LC line, aimed at helping clients on the path to becoming cognitive businesses and furthering IBM's commitment to open and collaborative innovation:

- IBM intends to add Open Compute Project-compliant systems to its Power Systems LC portfolio to support big data analytics and cognitive applications in the cloud. This is in addition to three other OpenPOWER Foundation members that recently announced plans for Open Compute Project-compliant, OpenPOWER systems: Mark III Systems, Penguin Computing and Stack Velocity.
- SUPERMICRO is currently developing two new POWER-based servers for IBM. The systems are based on the company's "Ultra" architecture and IBM intends to add them to the LC server line to add further design options. The two systems – a storage rich 2 socket, 2U design and a dense 2 socket, 1U design – will be POWER-based, GPU and CAPI acceleration enabled and fine-tuned for cloud and cognitive workloads.
- IBM, in collaboration with NVIDIA and Wistron, plans to release its second-generation OpenPOWER high performance computing server **[LINK TO SUMIT BLOG]**, which includes support for the NVIDIA® Tesla® Accelerated Computing platform. The server will leverage POWER8 processors connected directly to the new NVIDIA Tesla P100 GPU accelerators via embedded NVIDIA NVLink™ high-speed interconnect technology. Early systems will be available in Q4 2016. Additionally IBM and NVIDIA plan to create global acceleration labs to help developers and ISVs port applications on the POWER8 and NVIDIA NVLink™ based

platform.

About the OpenPOWER Foundation

The OpenPOWER Foundation is a global, open development membership organization formed to facilitate and inspire collaborative innovation on the POWER architecture. OpenPOWER members share expertise, investment and server-class intellectual property to develop solutions that serve the evolving needs of technology customers.

The OpenPOWER Foundation enables members to customize POWER CPU processors, system platforms, firmware and middleware software for optimization for their business and organizational needs. Member innovations delivered and under development include custom systems for large scale data centers, workload acceleration through GPU, FPGA or advanced I/O, and platform optimization for software appliances, or advanced hardware technology exploitation.

For further details visit www.openpowerfoundation.org.

###

OpenPOWER Foundation Reveals New Servers and Big Data Analytics Innovations

Foundation Membership Surpasses 200, Members Showcase More Than 50 Innovations at Summit

OPENPOWER SUMMIT, San Jose, Calif. – April 6, 2016 – The OpenPOWER Foundation today revealed more than 50 new infrastructure and software innovations, spanning the entire system stack, including systems, boards, cards and accelerators. Unveiled at the second annual OpenPOWER Summit, these new products build upon 30 OpenPOWER-based solutions already in the marketplace.

"*To meet the demands of today's data centers, businesses need open system design that provides greater flexibility and speed at a lower cost*," **said Calista Redmond, President of the OpenPOWER Foundation and Director of OpenPOWER Global Alliances, IBM**. "*The innovations introduced today demonstrate OpenPOWER members' commitment to building technology infrastructures that provide customers with more choice, allowing them to leverage increased data workloads and analytics to drive better business outcomes.*"

The OpenPOWER Foundation has rapidly grown to more than 200 businesses, organizations and individuals across 24 countries since it was formed two years ago. With the new innovations being announced today, the OpenPOWER Foundation continues to provide the technology and collaboration tools needed to deliver customized solutions and increased performance to customers, including hyperscale data centers and high performance computing. OpenPOWER innovations are built upon by a growing community of more than 2,300 ISVs supporting Linux on POWER applications.

The products revealed by OpenPOWER members today highlight:

· **New Servers for High Performance Computing and Cloud Deployments** – Foundation members introduced more than 10 new OpenPOWER servers, offering expanded services for high performance computing and server virtualization:

- Rackspace has announced that "Barreleye" has moved from the lab to the data center. Rackspace anticipates "Barreleye" will move into broader availability throughout the rest of the year, with the first applications on the Rackspace Public Cloud powered by OpenStack.
- IBM, in collaboration with NVIDIA and Wistron, plans to release its second-generation OpenPOWER high performance computing server, which includes support for the NVIDIA® Tesla® Accelerated Computing platform. The server will leverage POWER8 processors connected directly to the new NVIDIA Tesla P100 GPU accelerators via the NVIDIA NVLink™ high-speed interconnect technology. Early systems will be available in Q4 2016. Additionally, IBM and NVIDIA plan to create global acceleration labs **(link to Sumit blog)** to help developers and ISVs port applications on the POWER8 and NVIDIA NVLink-based platform.
- With planned availability in April, the TYAN GT75-BP012 is a 1U, POWER8-based server solution with the ppc64 architecture. The OpenPOWER-based platform offers exceptional capability for in-memory computing in a 1U implementation.

· **Expanded use of CAPI for Acceleration Technology** – Foundation members, including Bittware, IBM, Mellanox and Xilinx, unveiled more than a dozen new accelerator solutions based on the Coherent Accelerator Processor Interface (CAPI). Alpha Data also unveiled a Xilinx FPGA-based CAPI hardware card at the Summit. These new accelerator technologies leverage CAPI to provide performance, cost and power benefits when compared to application programs running on a core or custom acceleration implementation attached via non-coherent interfaces. This is a key differentiator in building infrastructure to accelerate computation of big data and analytics workloads on the POWER architecture.

· **A Continued Commitment to Genomics Research** – Following successful collaborations with LSU and the TransSMART Foundation, OpenPOWER Foundation members continue to develop new advancements for genomics research. Today, Edico Genome announced the DRAGEN Genomics Platform, a new appliance that enables ultra-rapid analysis of genomic data, reducing the time to analyze an entire genome from hours to just minutes, enabling healthcare providers to identify patients at higher risk for cancer before conditions worsen. DRAGEN's unprecedented speed is being leveraged to rapidly diagnose critically ill newborns, improve turnaround time for prenatal tests and quickly identify infectious disease outbreaks.

Developed in collaboration with Xilinx and IBM, the solution features Edico's DRAGEN processor, which is based on Xilinx's Virtex-7 980T FPGA, running on IBM Power Systems S822LC. The combination of the POWER CPU, high memory bandwidth and DRAGEN accelerated speed and high accuracy will allow clients, such as Rady Children's Institute for Genomic Medicine of San Diego, to leverage advanced analytics in genomics and life sciences.

For a complete list of member innovations announced today at OpenPOWER Summit, visit Link to Blog with Fact Sheet.

About OpenPOWER Foundation

The OpenPOWER Foundation is a global, open development membership organization formed to facilitate and inspire collaborative innovation on the POWER architecture. OpenPOWER members share expertise, investment and server-class intellectual property to develop solutions that serve the evolving needs of technology customers.

The OpenPOWER Foundation enables members to customize POWER CPU processors, system platforms, firmware and middleware software for optimization for their business and organizational needs. Member innovations delivered and under development include custom systems for large scale data centers, workload acceleration through GPU, FPGA or advanced I/O, and platform optimization for software appliances, or advanced hardware technology exploitation.

For further details visit www.openpowerfoundation.org.

###

U-Michigan, IBM collaborate on data-centric high performance computing system

San Jose, Calif. – April 6, 2016 – The University of Michigan is collaborating with IBM to develop and deliver "data-centric" supercomputing systems designed to increase the pace of scientific discovery in fields as diverse as aircraft and rocket engine design, cardiovascular disease treatment, materials physics, climate modeling and cosmology.

The system is designed to enable high performance computing applications for physics to interact, in real time, with big data in order to improve scientists' ability to make quantitative predictions. IBM's systems use a GPU-accelerated, data-centric approach, integrating massive datasets seamlessly with high performance computing power, resulting in new predictive simulation techniques that promise to expand the limits of scientific knowledge.

The collaboration was announced today in San Jose at the second annual OpenPOWER Summit 2016. The OpenPOWER Foundation, which U-M recently joined, is an open, collaborative, technical community based on IBM's POWER architecture. Several other Foundation members contributed to the development of this new high performance computing system, which has the potential to reduce computing costs by accelerating statistical

inference and machine learning.

Working with IBM, U-M researchers have designed a computing resource called ConFlux to enable high performance computing clusters to communicate directly and at interactive speeds with data-intensive operations. Hosted at U-M, the project establishes a hardware and software ecosystem to enable large-scale data-driven modeling of complex physical problems, such as the performance of an aircraft engine, which consists of trillions of molecular interactions. ConFlux, funded by a grant from the National Science Foundation, aims to advance predictive modeling in several fields of computational science. IBM is providing servers and software solutions.

"*There is a pressing need for data-driven predictive modeling to help re-envision traditional computing models in our pursuit to bring forth groundbreaking research,*" **said Karthik Duraisamy, assistant professor in the U-M Department of Aerospace Engineering and director of U-M's Center for Data-driven Computational Physics**. "*The recent acceleration in computational power and measurement resolution has made possible the availability of extreme scale simulations and data sets. ConFlux allows us to bring together large scale scientific computing and machine learning for the first time to accomplish research that was previously impossible.*"

ConFlux meshes well with IBM's recent focus on data-centric computing systems.

"*Scientific research is now at the crossroads of big data and high performance computing,*" **said Sumit Gupta, vice president, high performance computing and data analytics, IBM**. "*The explosion of data requires systems and infrastructures based on POWER8 plus accelerators that can both stream and manage the data and quickly synthesize and make sense of data to enable faster insights.*"

U-M researchers understand the significance of IBM's shift to data-centric systems, **said Michael J. Henesey, vice president business development, data centric systems and innovation centers at IBM.**

"*They were enthusiastic about the application of this architecture to problems that are essential to the university and to the country*," **Henesey said**. "*We will stay close to U-M to help inform our future system designs.*"

Progress in a wide spectrum of fields ranging from medicine to transportation relies critically on the ability to gather, store, search and analyze big data and construct truly predictive models ofcomplex, multi-scale systems.

Advanced technologies like data-centric computing systems are at the forefront of tackling these big data challenges and advancing the pace of innovation. By moving computing power to where the data resides, organizations of all sizes can maximize performance and minimize latency in their systems, enabling them to gain deeper insights from research. These data-centric solutions are accelerated through open innovation and IBM's work with other members of the OpenPOWER Foundation.

The incorporation of OpenPOWER technologies into a modular integrated system will enable U-M to configure the systems for their specific needs. ConFlux incorporates IBM Power Systems LC servers, which were designed based on technologies and development efforts contributed by OpenPOWER Foundation members including Mellanox, NVIDIA, Tyan and Wistron. It is also powered by the latest additions to the NVIDIA® Tesla®

Accelerated Computing Platform: NVIDIA Tesla P100 GPU accelerators with the NVLink™ high-speed interconnect technology.

Additional data-centric solutions U-M is using include IBM Elastic Storage Server, IBM Spectrum Scale software (scale-out, parallel access network attached storage), and IBM Platform Computing software.

In an internal comparison test conducted by U-M, the POWER8 system significantly outperformed a competing architecture by providing low latency networks and a novel architecture that allows for the integrated use of central and graphics processing units.

As one of the first projects U-M will undertake with its advanced supercomputing system, researchers are working with NASA to use cognitive techniques to simulate turbulence around aircraft and rocket engines. They're combining large amounts of data from wind tunnel experiments and simulations to build computing models that are used to predict the aerodynamics around new configurations of an aircraft wing or engine. With ConFlux, U-M can more accurately model and study turbulence, helping to speed development of more efficient airplane designs. It will also improve weather forecasting, climate science and other fields that involve the flow of liquids or gases.

U-M is also studying cardiovascular disease for the National Institutes of Health. By combining noninvasive imaging such as results from MRI and CT scans with a physical model of blood flow, U-M hopes to help doctors estimate artery stiffness within an hour of a scan, serving as an early predictor of diseases such as hypertension.

Studies are also planned to better understand climate science such as how clouds interact with atmospheric circulation, the origins of the universe and stellar evolution, and predictions of the behavior of biologically inspired materials.

"*The ConFlux project aligns with U-M's comprehensive strategy of investment in research computing and data science across disciplines,*" **said Eric Michielssen, U-M's associate vice president for research computing**. "*For example, our $100 million Data Science Initiative is advancing faculty driven research in engineering and the social and health sciences by building connections between the worlds of Big Data and HPC. ConFlux epitomizes this forward-looking vision.*"

For more information:

U-M's Center for Data-Driven Computational Physics

IBM's vision for data-centric systems