



Research Report

The OpenPOWER Foundation

Introduction

Just over one year ago IBM announced the formation of the OpenPOWER Foundation, an organization dedicated to expanding the ecosystem that surrounds IBM's POWER8 (and successive generations) processors. More specifically, IBM announced that it would make processor specifications, firmware and related software available to the evolving OpenPOWER community – essentially making the POWER processor an *open standard* microprocessor.

From our perspective, the reasons why IBM opened its POWER architecture were three-fold: 1) the company had seen a steep decline in sales of its Power Systems (since addressed – IBM now offers a more flexible consumption model known as mobile partitioning to make it easier for its customers to purchase Power Systems capacity) and needed to find a way to accelerate sales; 2) the company recognized that a broader ecosystem would foster even greater innovation – resulting in new POWER-based systems designs that would also drive sales; and, 3) IBM recognized and responded to changing market dynamics relative to Moore's Law (silicon alone isn't going to provide the necessary cost/performance improvements in server designs any longer – hence, new aggressive systems designs that exploit other components and that use accelerators are needed).

Our initial reaction to the formation of the OpenPOWER Foundation was extremely positive. We like IBM's open architecture approach with its POWER architecture for several reasons:

- It creates strong competitive pressure on the Intel x86 processor and its associated ecosystem. This pressure should drive Intel and its partners to build a more competitive products – which is good for consumers and good for the industry;
- It broadens the POWER ecosystem. Until the OpenPOWER Foundation was formed, IBM was the controller of POWER-based server designs. Now, with an open ecosystem in place, IBM's partners can build their own POWER-based systems designs – placing control of future POWER-based systems designs in the hands of the open community. As a result of open collaboration, new specialized servers are coming to market. Open innovation is a good thing for the industry; and,
- The foundation opens new markets for POWER-based systems. For various reasons, some geographies have been resisting the purchase of U.S.-made computer hardware and software products. These geographies can now create and manufacture their own POWER-based systems – and market those solutions globally.

In this *Research Report*, *Clabby Analytics* takes a closer look at the progress of the OpenPOWER Foundation. We describe what it is and how it is organized; we share some insights from several foundation members; and we discuss future POWER-based system designs. From our perspective, with over fifty members in its first year (and more than one hundred companies expressing interest), the OpenPOWER Foundation is off to a great start. With new accelerator systems designs coming to market as a result of this alliance, *Clabby Analytics* expects the OpenPOWER Foundation to become a major force in the development of workload optimized Big Data/cloud systems designs of the future.

The OpenPOWER Foundation

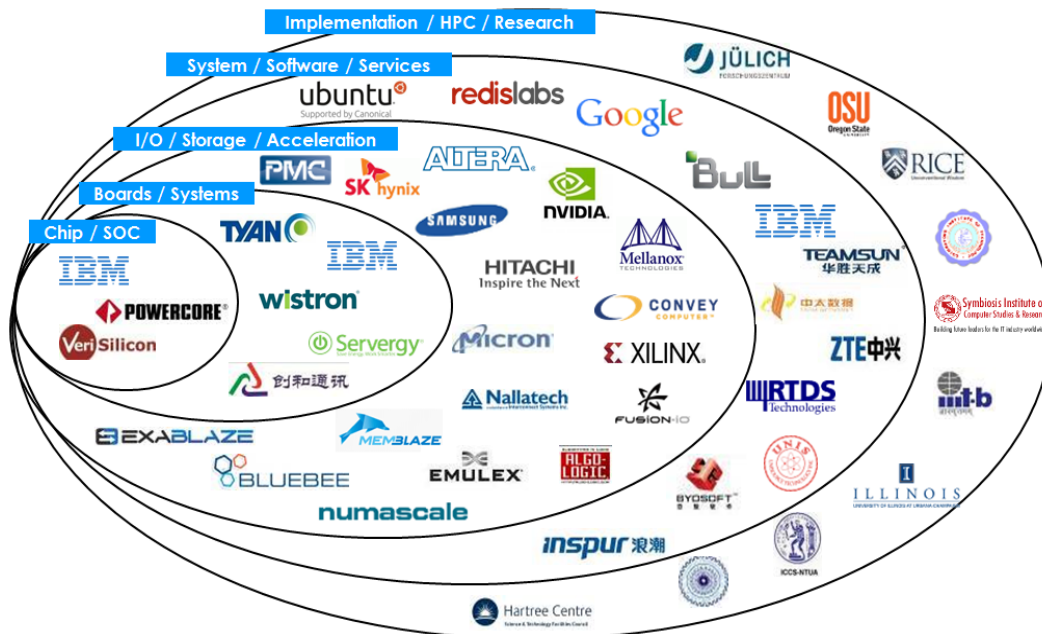
Overview: The OpenPOWER Foundation

The goal of the OpenPOWER Foundation is to foster innovation around the POWER architecture. Note: this foundation is not just about members communicating and innovating with IBM – the foundation is also about members creating innovative new designs with each other (independent of IBM). These new innovations will significantly increase the relevance of POWER in the IT marketplace – which in turn will drive the sales volumes that IBM is looking for.

A closer look at the current membership of the OpenPOWER foundation finds that there are now 59 members (as of October 1, 2014) – and, according to IBM, a hundred more vendors have expressed interest. Currently, member organizations fit into five processor/system design categories (illustrated in Figure 1). These categories are:

1. Chip makers/system-on-a-chip (SOC) makers;
2. Board makers/systems designers;
3. Input/output (I/O), storage, and acceleration product markers;
4. Systems, software, service providers; and,
5. Implementation, high-performance computing (HPC); and researchers.

Figure 1: OpenPOWER Foundation – Members and Specialties



Source: IBM Corporation – September, 2013

The list of accomplishments at the OpenPOWER Foundation over this past year include:

- *IBM* – Revealed new POWER8 Power Systems with CAPI (an interface described later in this report), published ABI specs and released 420,000 lines of firmware code;
- *Google* – Revealed single-socket reference board and contributed to firmware code release;
- *Tyan* – Revealed a single-socket reference board incorporating POWER processors;
- *NVIDIA* – Revealed GPU acceleration leveraging CAPI on POWER8 allowing clients to exploit GPUs for an 8x performance improvement;

The OpenPOWER Foundation

- *Mellanox* – Demonstrated network acceleration with a 10x improvement through high speed RDMA networking;
- *Xilinx* – Demonstrated KVS accelerations using a Xilinx CAPI attached FPGA accelerator;
- *Altera* – Demonstrated improvement in Monte Carol financial simulations run on an FPGA accelerator utilizing CAPI ;
- *Nallatech* – Revealed plans to integrate FPGA acceleration solutions;
- *Algo-Logic* – Revealed plans for application level solutions leveraging FPGA acceleration; and,
- *Inspur, Servergy, Convey Computer* and others have plans to build new servers and solutions based on OpenPOWER innovation.

It is worthy of note that 20% of the current members of the OpenPOWER Foundation are Chinese companies. And note that a major Chinese PC company (Lenovo) has just acquired IBM's System x (x86-based servers) business – a move that we project could make China the world's leading volume server maker in five to seven years. But, as noted above, other Chinese companies are showing interest in building specialized, even more powerful servers using IBM's POWER microprocessor as the CPU basis of their systems designs. It will be fascinating to see if the Chinese can build POWER-based servers that can meet Chinese market needs for powerful computers while addressing requirements for high availability, reliability and security. If new Chinese POWER-based systems can meet these design requirements, we can easily see Chinese manufacturers expanding to other markets throughout the world.

Member Commentaries: Why Join the OpenPOWER Foundation?

At a recent IBM Systems and Technology Group (STG) analyst briefing held in Stamford, Connecticut, *Clabby Analytics* attended a panel discussion hosted by Ken King, IBM's General Manager for OpenPOWER Alliances. Other members of this panel included Lixin Zhang - VP, Suzhou Powercore Technology Company; Albert Mu - GM, Tyan Computer; and John Zannos - VP, Cloud Channels & Alliances, Canonical.

When asked why each company had joined the OpenPOWER Foundation, Lixin Zhang of Suzhou Powercore Technology observed that his company joined the foundation for profitability reasons. Zhang described what it was like working with Intel, stating that “*Intel gets the meat, while we get the bones*”. IBM, on the other hand, offers its members a more liberal (and potentially more profitable) licensing agreement which closely resembles the ARM (advanced RISC machine) distribution model. Zhang also stated his belief that by the end of this decade China would surpass the US as the world's largest volume server manufacturer – implying that China has the expertise to build highly reliable, highly available, secure servers that will suit global needs. Zhang believes his company can build enterprise class servers based on IBM's POWER microprocessor technology.

Like Suzhou Powercore Technology, Albert Mu, general manager of Tyan Computer – a manufacturer of computer motherboards – also sees big opportunities in China and throughout the world for computer systems built with POWER microprocessors. Tyan currently builds motherboards with Intel and AMD processors – and is looking forward to applying its system design skills to POWER-based solutions.

Canonical's John Zannos, vice president of Cloud Channels and Alliances, sees the OpenPOWER Foundation as a means to broaden Ubuntu Linux marketshare. At present, the company has migrated 40,000 packages to Ubuntu POWER. This migration was accomplished over a five month period using the company's accelerated “juju” application testing/deployment environment.

The OpenPOWER Foundation

Zannos stated that “95% of Linux applications that are migrated from x86 require only a recompile and then they are operational”. Ubuntu Linux is on a six month release cycle (which is important to many ISVs), so the OpenPOWER Foundation gives Canonical/Ubuntu a time-to-market advantage against its primary competitors. As an example of how fast juju is, consider Zannos’ claim that it took four minutes to deploy IBM’s WebSphere middleware environment on Ubuntu – as Zannos put it: “that’s really fast!”

Innovation and the Arrival of Accelerator Architecture

The way we see it, the OpenPOWER Foundation is a technology incubator. Members invest research and development funds into projects that can exploit the power of the POWER microprocessor – and, as a result, the computing industry gets new systems designs that can process workloads faster than ever before. What we hope to see developed by the OpenPOWER Foundation over time is a bunch of new, highly-optimized *accelerator system* designs. These designs will use several different types of processors – and capitalize on the strengths of those processors – to deliver new high-performance systems.

Several examples of these new accelerator systems can be found in the Intel ecosystem today. For instance, last year we reported on three new systems designs that involved the use of Intel Xeon processors with other types of processors – resulting in accelerated workload processing:

- Our [report](#) on VelociData describes an accelerated system design that can analyze streams of structured and unstructured data at line speed using x86 CPUs (central processing units), graphical processing units (GPUs) and field programmable gate arrays (FPGAs) that accelerate large volume Big Data processing on varied data sets for complex workloads.
- Our [report](#) on IBM’s “The Now Factory” describes an acceleration system design that uses x86 processors and FPGAs to rapidly read and analyze Big Data, and is complimented by integrated infrastructure and applications software that can help communications service providers easily generate reports that provide deep network and customer insights; and,
- This [report](#) on system designs includes a discussion of IBM’s PureData for Analytics, another x86/FPGA design that has been designed and optimized to process complex analytics workloads.

The First Example of a POWER-based AcceleratorSystem

As part of its OpenPOWER Foundation commitment, IBM’s Power Systems organization has been working with fellow members of the foundation to build its own accelerator systems designs. One such design involves using NVIDIA GPUs to accelerate analytics, Big Data, and Java workloads.

This month IBM announced a new, highly-optimized accelerator Power Systems/NVIDIA system design aimed at executing pattern extraction analytics workloads faster than ever before (see Figure 2 – next page). In this particular design, NVIDIA GPUs are linked with IBM POWER processors over a Power System S824L PCIe bus. Work that can be executed more quickly on a GPU is assigned to the NVIDIA processor; work that can be executed more quickly on the POWER processor is assigned to it. The resulting system accelerated the processing of various technical computing workloads – and is especially effective at processing pattern extraction workloads.

The OpenPOWER Foundation

Figure 2 – IBM's New Accelerated Power Systems/NVIDIA Server

- ✓ Runs pattern extraction analytic workloads faster
- ✓ Provides new acceleration capability for analytics, big data, Java, and other technical computing workloads
- ✓ Delivers faster results and lower energy costs by accelerating processor intensive applications

Power System S824L

- Up to 24 POWER8 cores
- Up to 1 TB of memory
- Up to 2 NVIDIA K40 GPU Accelerators
- Ubuntu Linux running bare metal



Source: IBM Corporation – September, 2013

This new accelerator system design is a huge step in the right direction for IBM's Power Systems organization. This design accelerates specific technical workloads – and offers a much stronger alternative to Intel-based/GPU accelerator systems designs. However, we see this new system as only the tip of the iceberg in POWER-based systems designs. When IBM announced its POWER8 processor, the company also announced a new interface to that processor known as the CAPI interface (coherence attached processor interface). This interface allows many different types of devices to communicate directly with the POWER processor at very high speeds. We expect to see dozens upon dozens of new systems designs come to market through the OpenPOWER Foundation collaboration that will exploit the new CAPI interface to deliver superfast high-performance POWER-based accelerator system solutions to market.

Summary Observations

Many information technology (IT) executives believe that Intel-based servers offer the lowest cost computing solutions, and that Intel-based servers are well suited to execute almost every workload. At *Clabby Analytics*, we have written several reports that dispel these beliefs – showing that Intel-based servers are not the best solution for communications heavy, high input/output workloads (see this [report](#)), nor are they the best at processing compute/data intensive workloads (see this [report](#)). Intel Xeon processors are very good at executing several workloads (particularly those that involve processing light, fast threads such as supercomputing applications) – but they are not the most optimal choice for every workload.

By comparison, Power Systems are ideal for processing compute/data-intensive workloads (due to their ability to process 8 threads per core; due to the amount of cache memory that they can exploit; due to their high-speed, low latency bus, and more...). And now, with the introduction of the first accelerated Power System, Power Systems have become even more powerful when it comes to processing specific compute/data-intensive technical workloads.

So why, if Power Systems are so powerful, did IBM decide to open its POWER architecture to the marketplace? At last month's STG analyst briefing IBM stated that "no one company can, or should, control an innovation agenda for an entire industry. Opening the IBM POWER processor architecture draws deeper collaboration and more compelling innovation. It is becoming essential to allow multiple innovations at once, optimizing through the entire stack".

The OpenPOWER Foundation

But we think other factors also contributed to the formation of the OpenPOWER Foundation:

- It became apparent years ago in the computing industry that Moore's Law (the observation that steady improvements in innovation enabled the doubling of transistors on a chip about every 18 months – the means by which the industry constantly increased computing power) had just about reached the end of its path. There are physical limitations at the molecular level that are preventing silicon to shrink even further. Accordingly, new materials and approaches must be found in order to meet the never ending demand for more computing power. IBM was one of the first companies to recognize the decline of Moore's Law – and has been spending its research and development dollars finding ways to optimize other systems components (including building accelerators) while exploring the use of new materials in chip manufacture.
- IBM needed to encourage broader adoption of POWER architecture – and by opening-up this architecture to other vendors, IBM can expect to see more feet-on-the-street (more sales representatives making calls) driving broader distribution in multiple geographies throughout the world;
- IBM realized that by broadening the POWER ecosystem it could drive even more innovation around POWER architecture;
- IBM recognized that broader distribution of POWER architecture will also create new software/services revenue opportunities; and,
- We suspect that research and development (R&D) expenditure may have also played a role in IBM's decision to open-up POWER architecture. IBM does not have an open checkbook for Power Systems R&D. IBM has specific initiatives that it wants to drive with its Power Systems (around cloud, analytics, mobile and social or CAMS – as well as around cognitive computing) – but the POWER microprocessor can also process a wide variety of non-CAMs workloads. By opening-up POWER architecture to new channels of distribution, and by enabling other vendors to innovate using the POWER microprocessor, POWER-based solutions can be sold into market niches where IBM is not focused.

For all of these reasons, opening-up POWER architecture was a wise decision by IBM.

Parting Comment

At this juncture, a little longer than a year after the OpenPOWER Foundation was formed, we are starting to see the fruits of opening-up POWER architecture. IBM's new NVIDIA/Power System – which results from OpenPOWER Foundation collaboration between IBM and NVIDIA – represents just the tip of the iceberg when it comes to accelerator systems designs. Expect the OpenPOWER ecosystem to build traditional systems designs that will compete directly with IBM's own Power Systems – but also expect the new OpenPOWER ecosystem to build dozens-upon-dozens of new, extremely powerful accelerator system architectures that will consistently drive computing performance to new heights.

Clabby Analytics
<http://www.clabbyanalytics.com>
Telephone: 001 (207) 846-6662

© 2014 Clabby Analytics
All rights reserved
October, 2014

Clabby Analytics is an independent technology research and analysis organization. Unlike many other research firms, we advocate certain positions – and encourage our readers to find counter opinions – then balance both points-of-view in order to decide on a course of action. Other research and analysis conducted by Clabby Analytics can be found at: www.ClabbyAnalytics.com.