



Global Knowledge®

Expert Reference Series of White Papers

A Survey of IBM Power Systems Initiatives in Large- Scale Computing

A Survey of IBM/Power Initiatives in Large-Scale Computing Applications

Iain Campbell, UNIX/Linux Open Systems Architect, eLearning Specialist

Introduction

In a paper published fifty years ago, Gordon Moore proposed his now-famous theory that the number of components per integrated circuit would double every two years. Moore's Law, as it became known, tends to be used primarily to describe the increase in computing capacity of CPUs. In tandem with the growth in CPU capacity, Moore also accurately predicted the exponential growth in volumes of data manipulated by progressively faster CPUs, as well as the ability to perform increasingly more complex operations on these volumes of data.

In 2013 Google reported that the number of unique webpages in its search index stood at 30 trillion unique pages; the size of the index data itself weighed in at 1,000 terabytes. Yet, searching that index for "Moore's Law" will reliably return a result in well under a second. Manipulating data on such a scale and at such speed until recently has been restricted to a few very large players like Google, Yahoo, or Microsoft, using custom crafted and closely guarded proprietary solutions.

In recent years there has also been a great deal of work done by the open source sector on large-scale computing projects. Perhaps the best example of this is the Hadoop project. Originating as the Nutch project at the University of Washington in 2002, Hadoop is now the primary implementation of the Map/Reduce methodology used to break down large problems into discrete chunks, each of which can be solved in parallel in a compute cluster. This produces linear scalability of computing capacity and hence the problem with cluster size.

Today, state-of-the-art has come to a point where IBM is integrating its developmentally mature Power computing architecture with a variety of IBM proprietary and open source projects to bring to the market off-the-shelf solutions allowing for broader adoption of large-scale computing capabilities for all customers.

In this paper we will examine four such initiatives: the use of Linux on Power, the IBM Data Engine for Analytics, IBM Cloud Manager with OpenStack, and the OpenPower Foundation.

Linux on Power

As an operating system, Linux is most often associated with x86 hardware. This was certainly true when Linux was in its early growth years in the academic and hobbyist market. Linux was a good, free operating system that ran on inexpensive commodity x86 hardware. In the commercial marketplace, however, IBM has always been a strong supporter of Linux as an alternative operating system to closed source systems. In 1999 then IBM Vice President Sam Palmisano commissioned a study of Linux that resulted in CEO Lou Gerstner announcing that in IBM's eyes Linux would become "strategic to its software and server strategy." That announcement was followed up by the establishment that year of the IBM Linux Technology Center. The following year IBM publicly pledged to invest \$1 billion in Linux and the open source movement. Within the next two years IBM was running Linux on its core zSeries mainframe hardware, in 2007 IBM was a founding member of the Linux Foundation, and in 2011 IBM's cognitive computing engine, Watson, famously won "Jeopardy." Watson ran on Linux and was physically implemented on Power architecture.

By this point Linux had become a significant development platform for large-scale and cloud computing projects. Recognizing this, in 2012 IBM announced that in the Power7 product line they would start shipping Linux-only models, released at a price point competitive with enterprise x86 hardware, and designed to support Linux KVM virtualization out of the box. In 2013 IBM announced a second \$1 billion investment, this one directed specifically

at promoting Linux on the Power platform. The fourth Linux Technology Center also opened that year in Montpellier, France, adding to existing Linux centers in Beijing, New York City, and Austin, Texas.

Perhaps the most important Linux announcement in 2013, however, was the launch of the OpenPower Foundation, making key features of the Power processor architecture available under license to third-party hardware developers.

Keep in mind, too, that with the 2014 sale of its xSeries product line to Lenovo, x86 is now the competition for IBM, hence it is imperative for IBM that Power servers are given the support they need to compete with their x86-based competitors.

So how does this translate into tangible elements of support Linux on Power available today? Here is a list of some key components:

1. Linux-specific Power systems

IBM currently offers several Power8 server packages specifically targeting Linux applications. They include 2U and 4U packages offering up to 24 Power8 cores and 2Tb memory priced to compete with Xeon-based servers.

2. KVM support

Power Linux servers offer support for KVM out of the box. This is accomplished using OpenPower Abstraction Layer (OPAL), an OpenPower Foundation project. The intent of OPAL is to make a Power8 server look much like an x86 server in terms of its ability to boot and run a KVM-enabled Linux kernel quickly and with minimal overhead. Thus an OPAL-enabled Power8 server effectively becomes a slide-in replacement for an x86-based server in a KVM based virtualization infrastructure.

3. Linux distributions

IBM has always supported the flagship Linux distributions, Red Hat (RHEL7) and SuSE (SLES12), and has recently added Ubuntu (rel 14 LTS) to the list of supported ppc64le distributions.

4. Linux SDK

IBM offers a free, comprehensive SDK for Linux on Power. It includes key Linux tools such as oprofile, Prof, valgrind, and autotools, along with IBM tools such as FDPR and the IBM Java SDK all integrated into an Eclipse-based IDE. This IDE can run directly on Power hardware, or on an x86_64 client in conjunction with the included Power8 Functional Simulator, allowing development for Power in the x86_64 client without the physical server. The current version of the Linux SDK converges on little-endian support, eliminating ambiguity over endianness issues.

5. RAS and Power management features

The Power server platform has always incorporated an extensive suite of RAS functionality. Power consumption is now also a key issue, especially in large-scale deployments. The Power8 chip incorporates an on-chip power management micro controller capable of regulating power on a core-by-core basis. While RAS and power management are not specifically Linux-related issues, they speak to the large-scale server farm deployments associated with x86 deployment, and so are important to place P8 in the same racks as the x86 servers sit today.

It seems reasonable from this to conclude that IBM is taking Linux very seriously as a large part of its future. We mentioned Hadoop earlier as a prime example of an open source, large-scale computing project. We will look next at how IBM is pulling Linux and Hadoop together into the IBM/Power ecosystem to provide a turnkey big data offering.

Data Engine for Analytics

In the last few years the term “big data” has come into common use to describe the challenge of dealing with the volumes of raw digital data currently being generated. What does that mean for open source and Power? We will first establish a definition of big data means and then examine one solution IBM has built to address it.

Over time, data accumulates in both variety and quantity. In the past year or so discussions around storage solution design have routinely used the word *petabyte*, where a year ago the same discussion would likely have used the word *terabyte*. Here is a statistic to consider: the amount of data accumulated every forty-eight hours today is about equal to the sum total of all information in human history generated until the year 2003. A paper published by IBM in 2012 suggested that at that time some 2.5 exabytes of daily was generated worldwide every day.

Like mining for gold, there is valuable information hidden in all that data, but, like gold ore, that data has to be processed for the value to be extracted. This has always been true, but only in the recent past has the infrastructure existed that makes this extraction a practical possibility for the very large amounts of data systems are now called on to deal with. As it is not possible to build single machines capable of reading and processing petabytes of data, the only practicable way of doing it is to divide up the work and distribute it across multiple machines. To facilitate this, an infrastructure that distributes and manages large quantities of data across a set of machines is a necessity. It is also necessary to construct data mining code in such a way that it can run in parallel across the same set of machines; preferably in such a way that the portion of the code executing on any one of the machines operates to the greatest possible extent on the subset of the data resident locally on that machine. This is not a trivial task.

In the last decade, the increase in both data volumes and data processing capacity has fueled a great deal of research in the area of parallel compute clusters. The IBM Data Engine for Analytics offering is a turnkey infrastructure solution that represents the current state of that art. Let us examine the four key components of the IBM Data Engine:

1. The Power8 CPU

The heart of the data engine is the Power8 processor. Introduced in April 2014, the Power8 is the latest in the long line of Power CPUs dating back to 1990. A single Power8 socket executes eight hardware threads on each of its twelve cores, for a total of ninety-six concurrent threads of execution. Benchmark testing indicates that the P8 currently outperforms the competition—be it Intel or Sparc—by a comfortable margin.

Probably more important than the core compute power available, is the fact that IBM is opening up the Power architecture via the creation of the OpenPower Foundation, allowing third-party vendors to implement Power-based hardware devices. The licensing model in place is similar to the model successfully used by ARM to establish itself very effectively in the mobile and embedded-systems market. This move to open the architecture to licensees thus creates the possibility for an entirely non-IBM self-sustaining ecosystem to grow up around the architecture. We will examine this in more detail in the final section of this paper.

2. Red Hat Linux

In the big data world, Linux has emerged as the operating system of choice, and as we have seen earlier, IBM has fully embraced Linux. For the Data Engine offering the Red Hat Enterprise Linux (RHEL) distribution was chosen for the core operating system. Using Linux allows the data engine to leverage the substantial library of currently available, open source big data utilities, while retaining the option of including additional proprietary value added IBM tools.

3. InfoSphere BigInsights and InfoSphere Streams

The dominant big data software stack is the open source Apache Hadoop project. Originally written in 2005, Hadoop became an Apache project in 2011, and as such, is distributed under the Apache open source license. The core components of Hadoop include a distributed file system, a computing resources manager, and MapReduce, the heart of the Hadoop computing model. MapReduce is the key element of the stack that divides the task at hand across multiple nodes and then aggregates the answer at the end. A typical Hadoop installation includes a variety of other supporting components collectively known as the Hadoop Ecosystem.

For the data engine, the core software component is InfoSphere BigInsights, which packages a complete Hadoop Ecosystem (specifically, Hadoop plus Pig, Jaql, Hive, HBase, Flume, Lucene, Avro, ZooKeeper, and

Oozie) along with a suite of IBM-developed, value-added utilities. These include installation tools, a web console, a text processing engine enabling identification of items of interest in documents and messages, an Eclipse plugin to aid in developing custom text analytic functions, a spreadsheet-like data analysis tool, JDBC integration with Netezza and DB2, extensions to Hadoop's job scheduler, support for LDAP-based authentication mechanisms, performance improvements to processing of text-based compressed data, and adaptive runtime control for Jaql jobs.

Also included in the data engine software stack is InfoSphere Streams. Streams is designed to allow the development of parallel analysis code that's capable of dealing with live data streams, in contrast to a Hadoop cluster, in which work is generally batch-oriented and operates on static data. Streams is an IBM product founded on work done at the IBM Watson Research Center. Streams clusters are capable of processing millions of events per second in streams including audio, video, geophysical, financial, or medical data.

4. Platform Computing

Two key products from the Platform Computing portfolio are included in the data engine. Platform Cluster Manager provides self-provisioning and management functions, allowing rapid deployment of servers and networks within the data engine infrastructure, as well as support for multiple tenants and a user self-service portal.

Storage within the data engine is handled by the General Parallel File System (GPFS). Recently rebranded as IBM Spectrum Scale, GPFS was formerly known as Platform Elastic Storage. This highly scalable and reliable cluster file system originated from the Tiger Shark file system, a research project developed in the early 1990s at the IBM Almaden Research Center. GPFS uses a distributed storage model in which storage workload is distributed across multiple nodes in the cluster, each of which have direct access to a subset of the total physical storage in the cluster. In this way, I/O operations are parallelized not only over physical storage devices, but also at a storage server level, allowing for very high throughput combined with near linear scalability. Additions to recent releases of GPFS include a big data friendly tool to distribute data optimally across Hadoop style processing clusters (the GPFS File Placement Optimizer) and the ability to make optimal use of SSD storage within the cluster.

When these four components are brought together to build a data engine, the basic hardware building block is a Power8 S822L server. This is a 2U dual-socket machine providing twenty-four Power8 cores running at 3.3 GHz and having 256GB of memory. Currently, data engines can be built with between one and sixteen S822L compute nodes. At least one Hardware Management Console (HMC) is necessary to manage the virtualization environment and each data engine includes a single Power8 S821L cluster management server.

Storage is provided by GPFS implemented within one or more Elastic Storage Server (ESS) appliances. Each ESS consists of two S822L Power8 machines and up to four DCS3700 storage drawers, capable of serving just under 1Pb per appliance, including 3Tb of SSD storage. The various networking needs of the data engine are handled by three Ethernet switches, supporting a mixture of speeds ranging from 1Gb to 40Gb.

A complete data engine product is thus a turnkey solution capable of supporting multiple concurrent Hadoop, Streams or other data analytics tasks. It arrives onsite, according to IBM, "as a complete, pre-assembled infrastructure with preinstalled, tested software." Supplied with a variety of tooling to enable rapid deployment of new analytic tasks, it is designed to scale in linear fashion, and, if contributions from the OpenPower Foundation members materialize as hoped, it promises to have good potential for cost-effective customization and workload-specific optimization to meet future demands. It is a significant product package clearly intended to maintain IBM as a major player in the big data market segment.

Hadoop is a technique designed to deal with large volumes of data; the concept of cloud computing is another large-scale computing technique that addresses the challenges of implementing computing infrastructure on a large scale. Let us next examine how IBM is addressing the cloud.

Cloud Manager with OpenStack

The visual image of a cloud has long been used to denote an infrastructure component the broad function of which it is useful to know, but whose detailed implementation can be taken for granted. Starting in the early 2000s the term began to be used more specifically for a particular case of such a situation, a case in which the broad function was to provide a generic service, such as an instance of a virtual machine (termed Infrastructure as a Service, or IaaS), and instance of an operating system (Platform as a Service, or PaaS), or an instance of an application (Software as a Service, or SaaS). Initial cloud implementations were largely proprietary (e.g., Microsoft Azure, Amazon Web Services) and were typically built with a combination of new and pre-existing software and firmware components. IBM introduced such a product under the name SmartCloud in 2011. While these cloud implementations were, and still are, very functional, they lacked a broad, common, vendor-neutral standardized infrastructure model.

In 2010 Rackspace and NASA jointly launched an open source cloud computing platform project. Named OpenStack, this project is now under the control of the not-for-profit OpenStack Foundation. Supported by more than 200 members (including industry leaders such as Cisco, Intel, Hewlett-Packard, Red Hat, SuSE, VMware, and IBM), OpenStack is rapidly becoming a major player, if not *the* major player in the cloud infrastructure market. OpenStack is a cloud computing platform of choice for significant corporate web presences such as WebEx, PayPal, eBay, and Rackspace.

In October 2014 the announcement of the IBM Cloud Manager with OpenStack offering indicated that going forward IBM is also throwing its weight behind OpenStack as the platform of choice for IBM Cloud offerings. So what used to be IBM SmartCloud is now IBM Cloud Manager with OpenStack.

About OpenStack

OpenStack can be described as an Infrastructure as a Service (IaaS) stack with a modular architecture. The core modules of OpenStack are Nova (managing computing resources, interfacing with hypervisors such as KVM, Hyper-V and PowerVM), Cinder (managing block storage, interfacing with storage arrays), and Neutron (interfacing with networks). Other key modules deal with secondary issues such as managing and storing virtual appliance images (Glance) and data objects (Swift); a security and identity management service (Keystone); a database management service (Trove); a tool to manage Hadoop or Spark clusters (Sahara); and, a web administration interface (Horizon).

The OpenStack project is currently on a six-month release cycle with new modules being added and enhanced with each release. The current OpenStack version is code named Juno and was released this past October. For more information about OpenStack and the Juno release, a good starting point is to visit wiki.openstack.org.

About IBM Cloud Manager

Cloud Manager can be described as a toolkit that makes OpenStack easily accessible, usable, and manageable in both SME and enterprise environments. OpenStack is an engine, it handles the underlying manipulation of hardware such as servers, network switches, and array controllers. It provides some user and administration interfaces with more modules and functionality being added at each release, but there are still gaps to be filled and edges to be polished to provide a manageable turnkey cloud solution. Cloud Manager adds value in the form of a collection of functionality built on and around the core mechanics of OpenStack to provide a fully engineered cloud deployment.

Major functions implemented by elements of Cloud Manager include:

- **Packaging:** OpenStack is a large open-source project consisting of multiple interconnected components, each of which is under constant development. As revisions are released change management becomes an issue, both within the OpenStack core, and more so when other components necessarily added to the core to complete a cloud solution are taken into consideration. IBM addresses this by shipping the complete cloud solution in a self-consistent set of rpm packages, created and maintained by IBM. This way, the Cloud

Manager with OpenStack product can be installed as a single software entity with its own versioning and patch management. The customer does not need to track changes in the OpenStack project or deal with OpenStack patch management. The current version of Cloud Manager with OpenStack is v4.2, released in December 2014. For more information, search online for IBM Product Announcement 214-381.

- **Compute Platform Support:** The OS Nova compute module is not tied to any one compute platform but there are details to be implemented for each platform to be supported. The Cloud Manager product includes these details in the package so that Nova can talk to your hypervisor platform out of the box. Cloud Manager supports Intel platforms via your choice of Linux KVM, Hyper-V, or VMware. IBM Power is supported via either PowerVC or Linux KVM, and there is also support for System z via zLinux KVM.
- **Enterprise Management:** Cloud Manager adds self-service portal and administrative interfaces as well as tools to manage libraries of virtual images, allowing easy deployment, cloning, modification, and backup/restore of VM instances and data objects. Cloud Manager incorporates the Chef cloud automation framework with its concepts of recipes and cookbooks as a basis for a sophisticated data model incorporating the ability to define and manage projects structured around defined business processes. User accounts having differing tasks and responsibilities such as End User, Cloud Manager, Virtualization Administrator, and Platform Administrator allows for the workload of managing the entire cloud infrastructure to be distributed to multiple users at all levels of the infrastructure stack.
- **Platform Resource Scheduler:** When IBM acquired Platform Computing in 2012, the flagship product that came with the deal was Platform Load Sharing Facility (LSF). Rebranded as Platform Resource Scheduler (PRS), this widely known and well-respected application is included with Cloud Manager, providing an array of capacity and availability management features including intelligent VM placement, automated runtime optimization, CPU load balancing, and HA support.
- **IBM Development:** The Cloud Manager offering is presented as an enterprise-ready package blending the best of open source and IBM value-added technologies. Significant development effort has gone into integrating the various components to deliver a single functional product, which will be managed as a single coherent product going forward. This should equate to significant time saving and contribute to a reliable and manageable turnkey solution.
- **IBM support:** When it doesn't do what you want it to do you have a single number to call.

Linux KVM, Power8, and OpenStack

I have now referred to Linux KVM support twice, first in the context of Power server's purpose built for KVM support, and also in the preceding section when listing the computing platform support options offered by Cloud Manager where it was pointed out that Cloud Manager supports the Power architecture via Linux KVM as well as the traditional PowerVM hypervisor. This is significant when you consider that Nova, OpenStack's computing module, can interface directly with Linux KVM. In a Power-based cloud, this eliminates the need for the Hardware Management Console (HMC) and the PowerVM hypervisor. You can do the same thing on Intel-based hardware, running KVM x86 directly on Intel, in similar fashion eliminating the need for VMware or an equivalent Intel architecture virtualization layer. So, whether you opt for Power or Intel, either way the Cloud Manager solution reduces cost and complexity, and also places Power and Intel head to head on a level playing field in terms of evaluating CPU price/performance. Benchmark test results for Power8 are impressive, and IBM clearly feels that the value proposition of the Power8 might very well surprise some people. Time will tell.

Clearly the cloud is a key market for any company in the large-scale computing space, and this move to embrace OpenStack suggests that IBM sees this as the cloud platform of the future. IBM is directing significant support to the project; in 2014, IBM ranked second in contributions to OpenStack integrated projects and more than 300 IBM employees are working on OpenStack in various different ways. This strategy of coming alongside open source is consistent with the support IBM has shown for Linux KVM as a virtualization platform in previous releases of cloud offerings. Indeed, IBM is showcasing Linux on Power8 as a compelling value proposition combining the best of what the open-source ecosystem and IBM value-added technology has to offer.

Linux on Power, the IBM Data Engine, and CloudManager with OpenStack are all examples of software applied at various levels to address large-scale computing challenges. We will conclude this paper with an examination of the OpenPower Foundation, an IBM initiative to open the Power hardware platform architecture to third-party hardware developers targeting large-scale computing applications.

The OpenPower Foundation

For the first time in IBM history, IBM is opening one of its core technologies to third-party licensees via the creation of the OpenPower Foundation. The licensing model in place is similar to that successfully used by ARM to firmly establish itself in the mobile and embedded systems market. Following a collaborative development model, the intent is to foster the creation of a non-IBM ecosystem to grow up around the architecture. Current OpenPower members, such as Tyan and Suzhou, are working on Power8 motherboards and CPUs. This has the potential to bring Power8 technology to the market at competitive price points, giving customers a wider range of hardware providers for either components or turnkey systems beyond only IBM. In May 2014, Google, a founding member of the Foundation, announced a prototype Power8-based server motherboard. Described as a test vehicle, it is nonetheless interesting to see the operators of possibly the largest single large-scale-computing infrastructure in the world taking such an interest in the Power platform.

From a technical perspective, one of the key enablers of the OpenPower initiative is the fact that with Power8 IBM has replaced the proprietary on-chip GX bus of previous Power chips with the industry standard PCIe bus and designed an interface called the Coherent Attach Processor Interface (CAPI), which allows PCI devices to directly access virtual address space. This opens a path to the creation of third-party PCIe-based coprocessor devices able to operate as task-specific extensions to a Power8 CPU. Recently, Nvidia and IBM announced the availability of the Nvidia Tesla GPU installed in a Power8 system, and Nvidia recently contributed their NVLink fast interconnect technology to the OpenPower initiative, promising future Power8 systems even tighter integration of Nvidia GPU units and Power8 CPUs. Other vendors are working on general purpose field programmable co-processors (FPGAs) and CAPI-attached flash memory units. This ability to easily add task-specific hardware capacity can yield significant improvements to execution times for big data parallel processing jobs. This is a significant innovation, with the potential to substantially increase the already impressive processing capabilities of the Power8 platform.

The ability to implement hardware-based computational accelerators is another of the more significant projects arising from this partnership. Next we will examine this in more detail, as an example of the potential offered by this opening of the processor architecture.

Math Co-Processors

The idea of a co-processor to perform specific math computations is not a new one. From the beginning of the x86 architecture, for example, Intel included a companion x87 chip that could be implemented to handle floating point arithmetic calculations offloaded from the main x86 CPU. In the early days, one could buy a PC with or without what was called a math co-processor. By the early 1990s, however, with the introduction of the 80586 CPU, die sizes had become small enough to incorporate the functions formerly performed by the co-processor directly on the main chip. The day of the separate co-processor, at least in the x86 family, had come to an end.

IBM's Power architecture never implemented an external math co-processor in the same way Intel did. Instead the IBM architects focused on providing maximum arithmetic calculation horsepower within the core CPU package. This compute capacity was generally available to any thread running on the CPU; it was not compute-task specific and ran in series with all other necessary CPU functions. This was in contrast to the external co-processor, which had the potential to do computation specific calculations in parallel with the main CPU.

This architecture changed with the introduction of the Power7+ chip. The Power7+ was implemented with a 32nm feature size, down from the 45nm feature size of the Power7, but the Power7+ retained the same die size of the Power7, resulting in almost double the number of transistors available to the chip architects. The chip designers opted to make use of some of those extra transistors to target two specific math-intensive use cases. One was the calculation of cryptographic algorithms; and the other was the calculation required to compress and uncompress real-memory pages demanded by the new Active Memory Expansion (AME) functionality

implemented by the Power7 servers. The solution entailed implementation using an on-chip co-processor built and coded specifically for computing a set of common cryptographic algorithms and AME-specific memory compression and decompression algorithms.

The Power8 chip carried that design forward, adding additional on-chip accelerators for Hardware Transactional Memory (HTM), Virtual Memory Management, and Partition Mobility. What's interesting is that in addition to these on-chip accelerators, the Power8 adds a generic capability to support an x87-style external accelerator. This has significant implications for the future of the Power architecture as it opens the possibility for third-party vendors to provide co-processors to enhance the core capabilities of the Power8 chip into any specialized direction. Thus, Power8 becomes a platform upon which any number of specialized computing engines can be built.

Coherent Attached Processor Interface (CAPI)

A key component needed to support this external co-processor is CAPI. One of the challenges to overcome in implementing a co-processor is integrating it into the architecture of its host machine. Speed is critical; getting data to the co-processor and getting answers back as fast as possible is imperative. In the past, co-processor implementations used a device-driver model to do this, but that adds layers of protocol between the main system's memory address space and the data that is addressed by the co-processor. Ideally, the co-processor should be able to address the same memory space as the main CPU, allowing it to operate as a peer of and in parallel with the main CPU. This key memory access model is what CAPI provides, eliminating the device driver bottleneck.

Access to CAPI technology is being offered by IBM through the Open Power Foundation. While not an open source technology, CAPI and its associated technology is available to anyone willing to pay the license fee to become a member of the foundation.

The NVIDIA Tesla GPU

One of the early adopters of CAPI is NVIDIA, a leading developer of graphical processing units (GPUs). Originally developed to facilitate computing the large volumes of data needed by graphic-intensive applications such as CAD and gaming, GPUs are at heart mathematical computation engines, and by appropriate coding, they can perform any kind of mathematical calculation requested. Since 2007, NVIDIA has been doing just that—repurposing their industry-leading proprietary graphics processing technology toward general-purpose number crunching. Currently the NVIDIA Tesla K40 GPU can be ordered as a CAPI attached GPU for Power8 servers. To support the Tesla GPU, NVIDIA also supplies a programming model and associated instruction set (Compute Unified Device Architecture [CUDA]) to make it possible for developers to easily and effectively harness the power of the Tesla GPU and bring it to bear on their computations.

One of the areas of interest that IBM has recently targeted for a Tesla-based solution is the acceleration of Java-coded applications. The IBM-developed CUDA4J library provides application program interfaces (APIs) that allow Java programs to control and direct work to the Tesla engine using the normal Java memory model. Early experiments with Tesla-accelerated Java applications have yielded speed improvements approaching or exceeding an order of magnitude, and promise better to come.

Reaching the Summit

The Power+NVIDIA combination has drawn the attention of one of the biggest supercomputer customers around—the U.S. Department of Energy (DoE). Responsible for both the Oak Ridge and Lawrence Livermore laboratories, the DoE operates some of the largest supercomputers in the world and has done so for a long time. In late 2014 the DoE announced that the next-generation flagship computers commissioned for these labs would be based on the IBM Power9 + NVIDIA Volta technology combination. The largest of these machines, codenamed Summit, is due for delivery in 2017. Taking over from Titan, an Opteron+Tesla-based system currently ranked as the second most powerful supercomputer in the world, Summit will be more than five times more powerful.

Conclusion

In April 2015, on the occasion of the fiftieth anniversary of Moore's Law, Bradley McCredie, president of the OpenPower Foundation, described the work of the foundation as, "A powerful birthday gift to Moore's Law." As we begin to approach absolute physical limitations such as the size of an atom and the speed of light itself, it is not reasonable to expect that the capacity doublings of the past fifty years can continue. Moore himself, in a recent interview, stated that he felt his famous assertion would wind down within the next decade. The amount of data being produced, however, does not seem to show any sign of slowing; if anything, it may in fact be exceeding Moore's Law's growth rates. As these absolute limitations on raw hardware capacity begin to be felt, the importance of looking at all of the aspects of a computing infrastructure in order to tackle the future's truly large-scale computing demands, is imperative. The work described in this paper is representative of some of the major initiative IBM is currently pursuing to meet this challenge.

Learn More

Learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge through training.

[IBM Platform Elastic Storage \(ES\) System Administration Basic 4.1 for Linux](#)

[Implementing Cloud Manager with OpenStack on Power Systems](#)

[InfoSphere BigInsights Foundation](#)

[Virtualizing Linux on Power with PowerKVM](#)

Visit www.globalknowledge.com or call **1-800-COURSES (1-800-268-7737)** to speak with a Global Knowledge training advisor.

About the Author

Iain Campbell is a mechanical engineer by profession. While responsible for managing several production automation labs at Ryerson Polytechnic University in Toronto, he became distracted by UNIX operating systems. His first experience of AIX was a PC RT used in the lab as a machine cell controller. Iain has been teaching and consulting in AIX and Linux since 1997. He is the author of *Reliable Linux* (Wiley, New York, 2002), as well as several technical papers and curriculum material. He holds LPI and Novell certifications in Linux administration, and is an IBM certified AIX Specialist.