High Value Insights with Big Data Analytics on IBM Power Systems

Sponsored by IBM

Srini Chari, Ph.D., MBA mailto:chari@cabotpartners.com
October, 2014

Cabot Partners Group, Inc. 100 Woodcrest Lane, Danbury CT 06810, www.cabotpartners.com

Executive Summary

The increasing use of smart phones and social media is a reality across many industries today. It is not just where and how business is conducted that is changing, but the speed and scope of the business decision-making process is also transforming because of Big Data Analytics (BDA). For this, the Hadoop infrastructure is increasingly being used.

Organizations have a lot of internal data in the back office (Systems of Records) and are gathering a lot more through stakeholder interactions (Systems of Engagements). All this voluminous data can be static or dynamic and exist in a variety of forms: structured (in existing enterprise IT systems like customer relationship management (CRM), inventory, and billing, etc.), unstructured (audio, video, social media, email, chats, etc.). Big Data Analytics enables organizations to frame strategic business questions and combine Systems of Records with Systems of Engagements to produce new High Value Systems of Insights.

Businesses are investing in Big Data Analytics to improve customer experience and loyalty, discover new revenue opportunities, detect fraud and breaches, improve patient outcomes, mitigate financial risks, and more. Likewise, BDA helps governments respond faster to emergencies, analyze terrorist threats better and more accurately predict the weather – all of which are vital for national security, public safety and the environment. The economic value of Big Data Analytics is immense.

But this requires organizations to align strategy, culture, business processes and information technology (IT) to continually produce high value insights. However, enterprises have limited budgets for major infrastructure investments. They need a cost-effective, high-performance, reliable and agile IT infrastructure to leverage their unique data assets to deliver the best possible business outcomes. Solutions based on the POWER8 processors deliver this.

Organizations of all sizes should actively consider the IBM Big Data Analytics portfolio of solutions optimized for Hadoop workloads. These solutions are anchored on the POWER8 processor that provides a cost-effective, open, scalable and innovative platform for Big Data Analytics.

Clients who invest in IBM Power Systems for Big Data Analytics could lower the total cost of ownership (TCO) with fewer more reliable servers compared to x86 alternatives (see client examples). But more importantly, these customers will also benefit from the high value delivered by the growing open ecosystem of IBM Partners (OpenPOWER Foundation) and game-changing innovations such as Coherence Attach Processor Interface (CAPI).

Cabot Partners
Optimizing Business Value

Systems of Insights: The Big Opportunity with Big Data Analytics

Not just Big Data but also Diverse Data

In a world where data is exploding in velocity, variety, and volume, it is useful to define 'Big Data'. One good definition, from McKinsey[1], is 'Big Data' are datasets that are beyond the ability of typical database software tools to capture, store, manage, and analyze. This definition is intentionally subjective — as technology advances, the size of datasets that qualify as Big Data also increases. Today, Big Data in many sectors range from a few dozen terabytes (TB) to multiple petabytes (PB). But it's not just Big Data but also Diverse Data.

Increasingly, companies use unstructured data to interact with stakeholders

Data can be structured, semi-structured, unstructured or multi-structured.[2] Structured data is data that has been clearly formed, formatted, modeled, and organized so that it is easy to work with and manage (e.g., relational databases and spreadsheets), typically found in *Systems of Records.* Unstructured data covers most of the world's information but does not fit into the existing databases for structured data. Further, unstructured data consists of language-based data (e.g., e-mails, Twitter messages, books) as well as non-language based data (e.g., images, slides, audios, videos). Semi-structured data follows certain implied structures of structured data but is not standard enough to meet the criteria needed for the types of management and automation normally applied to structured data — for example, some fields may be missing or contain information that cannot be easily described in a database system. Examples of semi-structured data would be traffic signs posted along highways or web pages. Multi-structured data refers to combinations of two or more types of data structures. Increasingly, organizations interact with their stakeholders through *Systems of Engagements* that involve some form of unstructured data.

High value insights requires integration and analysis of structured and unstructured data

An estimated 85% of data is unstructured and originates from sources such as audio, documents, emails, images, RFID, social media, video, web logs, and so on. Obtaining *Systems of Insights* from Systems of Records and Engagements is a challenge for many organizations. But it is also a game-changing business opportunity in order to deliver exceptional customer experience, enhance marketing effectiveness, increase operational efficiencies, reduce financial risks, and improve product quality and reliability … etc.

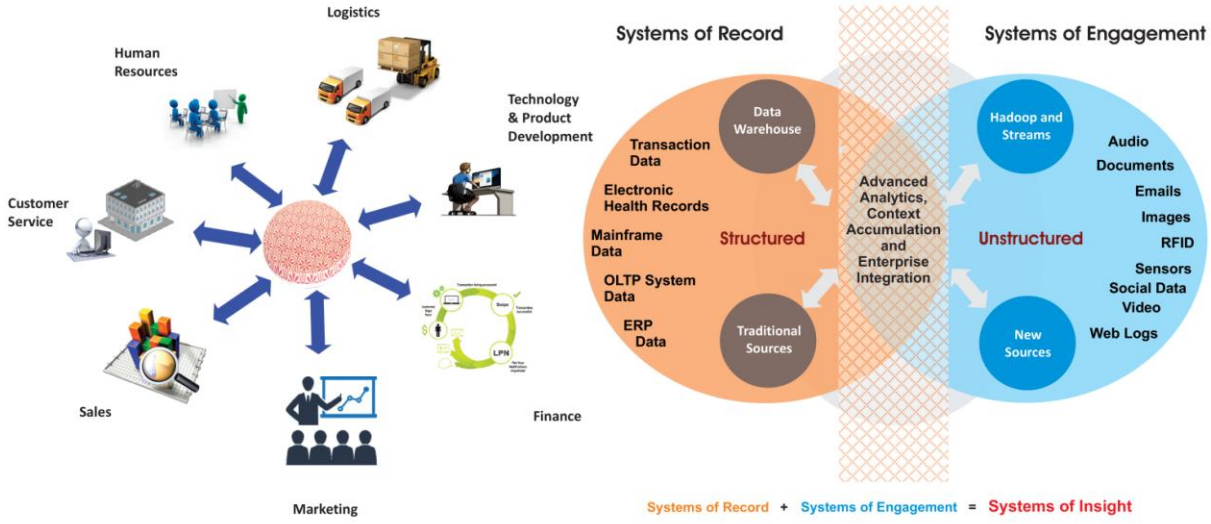

Figure 1: High Value Insights from Enterprise Integration of Structured and Unstructured Data

[1] Big Data: The next frontier for innovation, competition, and productivity, McKinsey Global Institute, June 2011.
[2] Big Data, Bigger Opportunities, Jean Yan, President Management Council Inter-agency Rotation Program, Cohort 2, April 2013.

Big Data Analytics enables iterative critical thinking processes for high value insights across many industries

Across many industry verticals, many Big Data Analytics high value use cases with iterative critical thinking processes are emerging. These allow organizations to frame strategic business questions and combine Systems of Records with Systems of Engagements to produce new *High Value Systems of Insights.*

McKinsey estimates that if the US health care system were to use Big Data Analytics creatively and effectively to drive efficiency and quality, the potential value from data in this sector could be worth more than $300 billion every year, two-thirds of which would be in the form of reducing national health care expenditures by about 8 percent. In the private sector, a retailer leveraging Big Data to the fullest extent has the potential to increase operating margin by over 60 percent. In the developed economies of Europe, government administration could save more than €100 billion ($149 billion) in operational efficiency improvements alone by using Big Data Analytics. This estimate does not include Big Data levers that could further reduce fraud, errors, and tax gaps.

But this requires organizations to better align strategy, culture and ***technology***[3] to get high value insights from Big Data Analytics. In fact, a recent IBM CEO Study indicates that CEOs consider technology factors their most important strategic issue – ahead of market factors.[4]

Enterprise-grade Information Technology Infrastructure is Crucial

IT is a key enabler and a primary CEO consideration

Traditional approaches to offline analysis or business intelligence with siloed data marts are limiting. They cannot keep up with the volume, variety and velocity of data that enterprises deal with today, as a result of next generation network roll-outs, proliferation of smart phones, rise of mobile/social media and the deluge of other unstructured data. Technical obstacles include slow data loading and querying, large network latencies, low system reliability and utilization, and the costs and complexities of managing distributed infrastructure.

To deal with these challenges and at the same time deliver new, revenue generating products/services without letting costs go out of control, many companies have begun to implement a high-performance, scalable and agile information foundation to support both real-time and large-scale analytics capabilities. These include the use of emerging open source technologies such as Hadoop and various NoSQL offerings to reduce the processing time for the growing volumes of data, especially in distributed computing environments. But, they also need robust Reliability-Availability-Serviceability (RAS), security and governance processes, normally found in enterprise-grade IT solutions such as the IBM portfolio of Big Data Analytics solutions anchored on the IBM **POWER8** processor. Key solutions in this family include the **IBM Solution for Hadoop – Power Systems Edition** optimized for Big Data workloads and the **IBM Data Engine for Analytics – Power Systems Edition** that supports and integrates multiple Big Data Analytics workloads efficiently.

Why POWER8 for Big Data Analytics

IBM solutions anchored on POWER8 for enterprise-class Big Data Analytics

With a data-centric design, IBM Power Systems offer a tightly-integrated and performance-optimized infrastructure for Big Data workloads. In addition, with the OpenPOWER initiative, IBM is building an ecosystem of partners. This is crucial to support Hadoop and other NoSQL open source Big Data Analytics infrastructures.

[3] IBM Institute for Business Value 2013 Big Data & Analytics Study, http://www.ibm.com/services/us/gbs/thoughtleadership/ninelevers/
[4] IBM Global CEO Study, 2012, http://www.ibm.com/services/us/en/c-suite/ceostudy2012/

High-performance POWER8 processors support massive threads and bandwidth with excellent RAS

Massive Threads and Bandwidth: The POWER8 is a massively multithreaded processor. Each core is capable of handling eight hardware threads simultaneously for a total of 96 threads executed simultaneously on a 12-core chip. The processor makes use of very large amounts of on- and off-chip eDRAM caches, and on-chip memory controllers enable very high bandwidth to memory and system I/O.

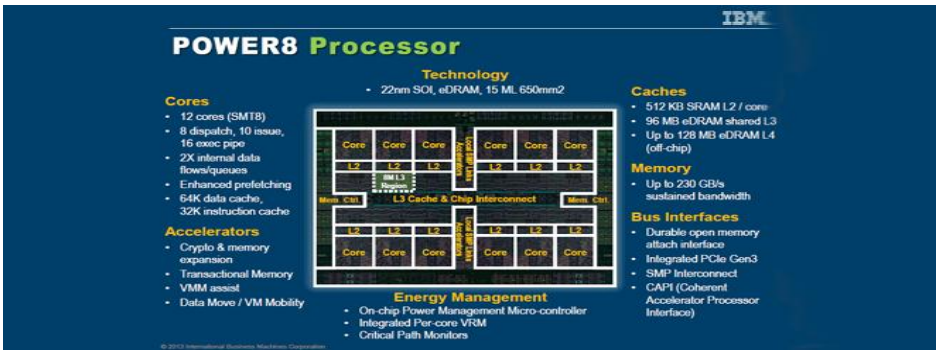

Figure 2: POWER8 Processor Details

Capable of analyzing Big Data workloads with IBM DB2 with BLU Acceleration 82 times faster than commodity x86 systems

High Performance Processor: POWER8 is capable of clock speeds around 4.5GHz, with a Thermal Design Power (TDP) in the neighborhood of 250 watts. At this speed, the POWER8 CPU will be around 60% faster than the POWER7 in single-threaded applications, and more than 2 times faster in multithreaded tasks. In certain cases, IBM says POWER8 is capable of analyzing Big Data workloads with IBM DB2 with BLU Acceleration 82 times faster than commodity x86 systems with a traditional row-based database with the same amount of RAM and same number of cores. In fact, IBM clients have reported seeing 50 times to over 1000 times faster analytic queries with BLU Acceleration on Power Systems.[5]

Excellent RAS and Virtualization: Recent studies[6,7] across a range of enterprises indicate that IBM Power systems perform better than x86 systems in RAS, performance, TCO, security and overall satisfaction. These results reinforce the fundamental advantages of the POWER8 architecture, including the ability to support high levels of concurrent error detection, fault isolation, recovery, and availability.

Advanced virtualization with PowerVM can be used to consolidate many different workloads/tasks including production databases, applications, high availability servers, backup/recovery servers and other servers onto a single, smaller set of servers. It is much easier to optimize parallel tasks found in many Big Data Analytics applications on fewer servers than on the many servers x86 solutions undoubtedly require.

CAPI is a game-changer

Coherence Attach Processor Interface (CAPI) – a Performance Game Changer: Beyond raw SPECint and SPECfp performance, POWER8 also introduces CAPI, a direct link into the CPU, allowing peripherals and coprocessors to communicate directly with the CPU, substantially bypassing operating system and driver overheads. While Intel's bus protocols and interfaces are proprietary, IBM has developed CAPI to be open to third party vendors and even offers design enablement kits. In the case of flash memory attached via CAPI, the overhead is reduced by a factor of 24:1. More importantly though, CAPI can be used to attach

[5] https://www14.software.ibm.com/webapp/iwm/web/signup.do?source=sw-infomgt&S_PKG=ov20419
[6] Edison Group, " Better Performance, Lower Costs The Advantages of IBM PowerLinux 7R2 with PowerVM versus HP DL380p G8 with vSphere 5.1, http://public.dhe.ibm.com/common/ssi/ecm/en/pol03161usen/POL03161USEN.PDF
[7] Solitaire Interglobal, "Power Boost Your Big Data Analytics Strategy", http://www.ibm.com/systems/power/solutions/assets/bigdata-analytics.html

CAPI-attached coprocessors can make a POWER8 system 1,000 times faster than a comparable x86 system

coprocessors — GPUs, FPGAs — directly to the POWER8 CPU for significant workload-specific performance boosts. These CAPI-attached coprocessors can make a POWER8 system 1,000 times faster than a comparable x86 system.

CAPI is an innovative approach to heterogeneous computing using accelerators. This is important because scale-out data centers of the future are likely to be workload and application specific. This means that servers would be dedicated to different workloads using specialty hardware that is optimized for that specific task. In other words, the server environment would not be homogeneous. With CAPI and OpenPOWER (see below), IBM has created an environment where specialty accelerators from NVIDIA, Altera, VeriSilicon, and other partners can be used to accelerate workloads far beyond what can be done today, and this can be done in an architecture that is easier to program than previous architectures.

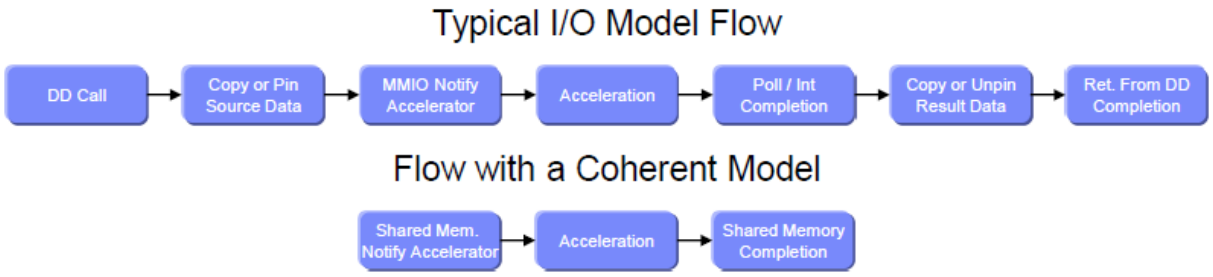

Figure 3: The CAPI Advantage

CAPI provides ultra low-latency large shared memory with traditional programming model to accelerate Analytics

The advantages of CAPI over a traditional I/O attachment include:

- Virtual Addressing and Data Caching
  - Large shared memory. Invaluable for data-intensive applications.
  - Lower latency for highly referenced data.
- Easier, more natural programming model
  - Traditional thread level programming avoids more complex distributed memory approaches.
  - Long latency of I/O typically requires expensive restructuring of application.
- Enables applications not possible with traditional I/O
  - For example, pointer chasing.

OpenPOWER Foundation opens up the POWER Architecture to a growing ecosystem to spur innovation

OpenPOWER Foundation Creates an Open Ecosystem for Innovation: The other big announcement with the POWER8 launch was IBM's reiteration of its commitment to the OpenPOWER Foundation (http://openpowerfoundation.org/). IBM has already invested over $2.4 billion in POWER8 technology over three years of development and is looking to the OpenPOWER Foundation to use that technology to bring additional innovations to the platform.

The OpenPOWER Foundation is a collaboration around Power Architecture-based products initiated by IBM and announced in August 2013. IBM is opening up technology surrounding Power Architecture offerings, such as processor specifications, firmware and software. IBM offers this on a liberal license basis and will be using a collaborative development model with partners. The goal is to enable the server vendor ecosystem to build their own customized server, networking and storage hardware for future data centers and cloud computing.

System stack innovations key to maintain cost / performance trajectory and accelerate performance real-world applications

The Foundation – representing over 55 global technology providers and growing – was founded by IBM, Google, NVIDIA, Mellanox and Tyan. The group has announced an innovation roadmap detailing planned contributions from several of its members, with IBM Power Systems as the first servers to exploit OpenPOWER technology. The underlying belief driving the OpenPOWER Foundation is that microprocessors alone can no longer drive cost/performance improvements. System stack innovations are required to improve cost/performance.

There are already several real world examples of innovations and performance enhancements resulting from the OpenPOWER Foundation and these span the business spectrum ranging from Monte Carlo financial risk modeling, Big Data and Java acceleration, NoSQL acceleration, Key Value Store (KVS) acceleration and so on.

Drives community innovation with support of key open standards

Leveraging open standards, IBM Power Systems provide developers with tools tuned for a platform that boosts productivity and performance by removing constraints imposed by commodity architectures. With an architecture that supports open standards such as Linux, OpenStack, KVM, and others—and with community innovation driven through Power Systems Linux Centers, Watson & Power development clouds, and the OpenPOWER Foundation—Power Systems signals the next generation of applications and infrastructure and promises significant acceleration in compute and data-intensive tasks.

The IBM Solution for Hadoop – Power Systems Edition

IBM Solution for Hadoop - Power Systems Edition provides an expertly-designed, tightly-integrated and performance-optimized reference architecture for Big Data workloads. Unlike traditional x86 infrastructures, this IBM Solution for Hadoop can be tailored specifically to meet Big Data workloads by using a simple building block approach to match the appropriate mix of CPU and storage to application requirements. This approach optimizes the infrastructure for Big Data workloads at a price that is competitive with x86 solutions. This architecture defines the following[8]:

Optimized, fully-supported cluster - customized for Hadoop to deliver the best performance and/or cost / performance

- **Complete cluster.** A comprehensive, tightly-integrated cluster designed for ease of procurement, deployment, and operation. This reference architecture describes all required components for Big Data applications, including servers, network, storage, operating system, management software, Hadoop compatible software, and runtime libraries. Customers have the flexibility to deploy storage in either a traditional Hadoop-oriented distributed topology, or in a shared disk topology that can provide more economical use of storage by avoiding the need for block-level replication.
- **Application optimized configuration.** The configuration of the cluster is designed to optimize application performance and reduce total cost of ownership. The cluster is integrated with IBM Platform Cluster Manager and IBM InfoSphere BigInsights Enterprise Edition which has IBM Platform Symphony MapReduce and IBM General Parallel File System – File Placement Optimizer (GPFS - FPO) as embedded components. This optimized configuration enables users to become productive more quickly.
- **Advanced technology for performance and robustness.** The underlying hardware and software components in the cluster are customizable to allow the best performance or the best price/performance ratio.

[8]IBM Solution for Hadoop - Power Systems Edition, an IBM Reference Architecture, June 2014.

Ability to dynamically adapt compute and storage is ideal for multiple Hadoop workloads

- **Commercial, solution-level support.** The configurations are validated and supported for all hardware and software components.

Because the architecture is optimized for general Big Data applications, and because most software deployments are fully automated, a cluster can be deployed in a day or two. Users are able to take advantage of the performance and capability that the cluster delivers with minimal training.

The IBM Data Engine for Analytics - Power Systems Edition

Typical Hadoop solutions use a storage-rich, server-based scale out solution. The challenge with such an approach is that users tend to lose control of storage since it is lumped with compute servers and there is no separate storage capital planning capability. Further, there are usually no backup, archiving, disaster recovery facilities, and storage security controls and audits are not rigorous. This can lead to heavy rebuild penalties on disk failures, consume network resources, and reduce application performance.

Elastic Storage based on IBM GPFS with a high performance interconnect alleviates many storage challenges with traditional Hadoop

On the other hand, multiple replicas for storage are expensive and superfluous if the workload doesn't need several tasks simultaneously accessing the same data over and over (mostly reading). This is also complicated by the fact that Hadoop solutions cannot typically share resources with non-Hadoop workloads and cannot reuse existing infrastructure. All this leads to a non-optimized configuration.

The IBM Data Engine for Analytics is designed to address these challenges by connecting POWER8-based compute servers to the Elastic Storage Server based on IBM General Parallel File System (GPFS) with a high performance interconnect. IBM GPFS is a high-performance clustered file system that can be deployed in shared-disk or shared-nothing distributed parallel modes and is commonly used in both enterprise computing and supercomputing environments. GPFS provides concurrent high-speed file access to applications executing on multiple nodes of clusters.

The IBM GPFS-based Elastic Storage Server used by the IBM Data Engine for Analytics offers shared storage at lower cost than SAN and includes RAS and management functionality.

IBM GPFS Elastic Storage Server reduces replication overhead, improves RAS and provides centralized common storage

There many advantages of using the Elastic Storage Server:

- Reduced replication overhead and less stress on network with elimination of replicas
- Faster rebuilds for disk failures with less network degradation and reduced impact on application performance
- Centralized storage management, better backup, archiving, disaster recovery, security controls and auditing
- Common infrastructure for different workloads, easier to build workflows that include Hadoop and non-Hadoop workloads, and solution level configuration, ordering, fulfillment, deployment, and support.

Fully integrates innovative capabilities of GPFS Elastic Storage Server, Platform Computing and networking components with Linux on Power Scale-Out Systems

Figure 4: Component Architecture of IBM Data Engine for Analytics – Power Systems Edition

POWER8 delivers 2X better performance versus x86 cores

According to IBM, the IBM Data Engine for Analytics offers several advantages over standard x86 Hadoop solution configurations:

POWER8 Advantage

- 4X more hardware threads per core. Better Simultaneous Multithreading (SMT) performance than Hyper-threading (HT)
- 30% more memory capacity
- 170% more memory bandwidth
- Built in hardware encryption/decryption engine
- 2X better performance versus x86 cores.

Platform Symphony provides a faster MapReduce engine with better resource utilization through multi-tenancy

Platform Symphony Differentiation

- Faster MapReduce engine
- Better utilization of resources that minimizes infrastructure duplication and reduces cost
- Allows sharing of infrastructure with non-Hadoop workloads
- Multitenant sharing of infrastructure
- Guaranteed SLA for tenants
- Journaling makes sessions and jobs recoverable – so, in case of failures, long-running jobs can resume without delays; saving time and money
- Security isolation between tenants
- Allow tenants to expand to use available infrastructure
- Low latency and faster throughput via User Space Sockets.

GPFS speeds up ingestion and sharing of data and simplifies integration of Hadoop and non-Hadoop workloads

GPFS Value

- POSIX file system makes it easier to build workflows that include Hadoop and non-Hadoop workloads. Easier and faster sharing and ingestion of data.
- Faster file system performance that translates to workload acceleration through native exploitation of RDMA capable networks.
- GPFS ACLs (access controls) allow better security control of data between multiple tenants in a shared infrastructure environment.
- GPFS Encryption and secure deletion functions adds more security.
- Distributed metadata server prevents a single point of failure and better performance than a single name node.
- Combines with other GPFS features like, Active File Management (AFM), Information Lifecycle Management (ILM), and Multi-cluster.
- Elastic Storage Server provides a faster and more resilient storage that also provides faster recovery from failed disks with minimal impact to application performance.

POWER8 Big Data Analytics Examples

Some examples showing the high value delivered by Power Systems for Big Data Analytics:

NCSU could run more concurrent queries faster on POWER8 - x86 systems didn't scale

North Carolina State University: The Center for Innovation Management Studies (CIMS) at the Poole College of Management in North Carolina State University (NCSU) reports that companies[9] ranging from startups to members of the Fortune 500 struggle with generating strategic inquiries from unstructured data. In case after case, the researchers found that the ability to make data-driven decisions requires close teamwork across the IT department to support the software and help gather and store the data, and among the statistical analysts to run the big data process.

The NCSU researchers also found that the right combination of software and server platforms is central to this process. They started the project using both x86 and Power-based systems but soon discarded the x86 based systems since they were unreliable and could not scale to Big Data environments. In their experience, the Power Systems server, with the POWER8 processor-based technology, was ideal for Big Data since it enabled them to run more concurrent queries in parallel faster, across multiple cores with more threads per core. It also has increased memory bandwidth and faster IO to ingest, move and access data faster.

Need more x86 servers for same performance and functionality

Bloor Research: Bloor Research, a leading UK based IT research firm and consultancy, recently compared the relative merits of IBM DB2 with BLU Acceleration running on Power Systems as compared to SAP HANA, Oracle Exadata and Microsoft SQL Server running on x86-based platforms.[10] The project was part of research to compare the performance capabilities of the leading Business Analytics platforms.

Bloor Research found that Power Systems offer substantial benefits for DB2 with BLU Acceleration because POWER8 has up to 12 cores per socket and 8 threads per core. This substantially increases the per socket parallelism to 96 for POWER8.

[9] The Business Case for Using Unstructured Text Analytics on IBM Power Systems for Critical Decision Making, Stephen Markham, Ph.D. and Michael Kowolenko, Ph.D., Poole College of Management, North Carolina State University.
[10] IBM DB2 with BLU Acceleration on Power Systems: how it compares, AnIn Comparison Paper by Bloor Research, Philip Howard, April 2014.

Bloor also pointed out that virtualization is built into Power Systems through microcode and so virtualizing Power servers does not typically have a significant impact on performance (provided workloads are balanced). That is not the case with x86 systems, since virtualization is not built-in and so you have to install software such as VMware, which can take up as much as 20% of system resources and thereby reduce performance.

Lower TCO on POWER8 solutions compared with equivalent x86 solutions

Also, Power Systems supports larger memory sizes. IBM offers capacity on demand for larger systems, and resiliency and security are more efficient in Power Systems. Solutions using x86 processors simply use redundancy – this means more servers, higher acquisition and maintenance costs and higher levels of complexity to manage.

Bloor concludes that though Power Systems appear to be more expensive than x86-based platforms, clients need more x86 servers to get the same performance and functionality as Power Systems. This drives up the cost for a x86 solution.

Conclusions

To get high value insights, organizations should actively consider Big Data Analytics solutions using IBM Power Systems with POWER8 and CAPI:

POWER8 delivers high value insights for Big Data Analytics on an open innovative platform

- All businesses could benefit from a lower total cost of ownership (TCO) with fewer servers (enabled by multi-tenant cluster virtualization which improves utilization). They also benefit from less storage and data bottlenecks because of consolidation, fewer redundant copies, novel compression algorithms and efficient data-aware scheduling.
- Many businesses already run mission-critical Data Warehouses and Business Intelligence Applications on Power Systems with excellent RAS and performance. Adding new analytics capabilities by mashing up this structured data with unstructured data is a lot easier to implement. It also protects current IT investments in people, processes, platforms and applications while providing a seamless and cost-effective path to scale to meet the demands of Hadoop and NoSQL workloads.

But most important, the POWER8 with CAPI provides an open innovative platform to enable the critical thinking needed for the iterative quest for **High Value Insights from Big Data.**