Advisory

Infrastructure Matters: POWER8 vs. Xeon x86

Executive Summary

This report compares IBM's new POWER8-based *scale-out* Power System to Intel E5 v2 x86-based *scale-out* systems. A follow-on report will compare *scale-up* POWER8-based Power Systems to Intel *scale-up* servers when those POWER8-based servers become available.

As the basis of our comparison we have chosen to compare IBM's one and two socket POWER8-based systems to Intel one and two socket Xeon E5-2692 v2-based (IvyBridge - EP) systems. We chose the Xeon E5 architecture because we see it as the most logical competitor to single-/dual-socket POWER8-based *scale-out* systems. (We see Intel's Xeon E7 [Ivy Bridge - EX] architecture positioned as the logical *scale-up* alternative given its support for larger quad- and octa-configurations).

From a microprocessor perspective, Intel's Xeon and IBM's POWER8 microprocessors are both "general purpose processors" – which means that each processor has been designed to execute serial, parallel and compute intensive tasks. But a closer look at each processor reveals significant differences in processor efficiency (particularly the number of threads that can be processed per clock cycle); in overall performance; in the amount of on-chip cache; and in the amount of memory that can be addressed.

As for system design, the biggest difference can be found in memory bandwidth speed.

> *POWER8-based uniprocessor/dual processor systems can execute more work more quickly than Intel E5 v2 Xeon-based servers. POWER8 can process four times as many threads as its E5 v2 competitors; the POWER8 clock speed is faster; and POWER8 can work on significantly more data in cache than E5 v2-based servers. Furthermore, POWER8 memory bandwidth is much faster than the Intel environment. What all of this means is that POWER8-based servers can <u>outperform</u> E5 v2-based servers (helping enterprises achieve results more quickly); and POWER8-based servers are also <u>more efficient</u> (meaning that enterprises will not need to buy as many servers to execute workloads. This will help enterprises save BIG MONEY by not having to purchase as many software licenses).*

In this *Research Report*, *Clabby Analytics* compares and contrasts Intel E5v2-based scale-out architecture with IBM POWER8-based scale-out architecture. Our key findings are:

1. Both microprocessor/server environments have been designed to process Web, file and print, email, database, vertical-specific applications, high performance computing and cloud workloads;

2. POWER8 processors are more efficient than Xeon processors;

3. Due to processing and bandwidth advantages, POWER8-based servers can deliver results more quickly;

4. POWER8-based servers are better suited for data-intensive environments; and,

5. When executing identical workloads, POWER8-based servers will cost less that E5 v2-based competitors (due to aggressive IBM pricing and numerous efficiency advantages).

Background

In June, 2013, Intel released the first member of its dual processor E5 v2 family – the Xeon E5-2692 v2 (Intel now offers 38 different E5 v2 processors that operate at varying speeds, with from four to fifteen cores per processor). These processors have been designed largely to serve the Windows and Linux marketplaces. There is a lot to like about this family of microprocessors as compared with the previous Xeon generation because these processors (codenamed "Ivy Bridge") offer more cores, more cache, faster speed, lower energy consumption, and new reliability/avail-ability/serviceability extensions). But probably the biggest improvement in Xeon v2 architecture is the amount of main memory that can now be addressed. With v2 architecture, x86 servers will someday be able to address up to 16TB of main memory in large, scale-up configurations.

In April, 2014, IBM introduced its new POWER8 microprocessor architecture as well as a new line of single- and dual-socket scale-out servers. IBM has focused these servers on serving both the *scale-out Unix* (AIX) marketplace as well as serving the *scale-out Linux* marketplace. (Previous generation POWER7 and POWER7+ servers could also run Linux – but now IBM has oriented its Power Systems to serve little endian as well as big endian applications – making it possible to more easily deploy Linux little endian applications on POWER architecture. Endians have to do with the byte order that systems read applications – and now POWER8 based systems can read Linux appli-cations in the same way that x86-servers can read them. This makes it easier to deploy thousands of Linux applications on POWER8 architecture – a situation that IBM hopes to leverage to grow its Linux on Power Systems user base).

> *Clabby Analytics is impressed with both architectures. We like the way Intel has finally addressed the memory limitations of scale-up x86 architectures (the subject of this [report]). But we are especially impressed with the processor efficiency, performance and bandwidth improvements offered with POWER8-based systems. With a faster clock speed; with the ability to process four times as many threads per cycle as x86 processors; with three times more on-chip cache; with four to six times the memory bandwidth – and with access to up to 80 TB of Flash using the newly introduced coherence attached processor interface (CAPI) – IBM has created a systems environment with POWER8 that has been designed to very significantly outperform Intel Xeon architecture.*
>
> *The way we see it, this new generation of POWER8-based servers has literally been "designed for data". We see the current generation of single and dual-socket systems as a true threat to the x86 dominance of the scale-out Linux marketplace. And, as larger and large configurations come to market, we expect to see a lot of POWER8-based servers configured for large database-in-memory processing – and given POWER8 performance advantages, these new in-memory servers will raise the performance bar for database processing in the future.*

The Primary Differentiators: Performance and Efficiency

As we compared Intel's Xeon E5-2692 v2 with IBM's POWER8 architecture from both a processor design and subsystem perspective, it became readily apparent that both chips were designed to process serial, parallel and data-intensive workloads – but it also became clear that IBM's POWER8 architecture was designed to deliver high-performance while operating far more efficiently than Intel's Xeon architecture.

To illustrate these points, consider the following:

- *Performance* – IBM's POWER8 is more than 25% faster per clock cycle than Intel's E5-2692 v2 (IBM operates at 4.15 GHz per clock cycle; Intel's E5-2692 v2 runs at 2.697 GHz). POWER8 offers three times more on-chip cache (data in cache can be read faster – leading to faster performance and faster results). Further, POWER8 can address almost

25% more main memory than the E5-2692 v2 – again placing more data closer to the processor where it can be read and acted-on more quickly. And POWER8 can receive data from memory four to six times faster than Xeon bus can. The combination of a faster processor with access to more cache and memory – with significantly faster memory bandwidth – make POWER8 processors more powerful than Xeon E5 architectures.

- *Efficiency* – POWER8 offers a 12 core processor configuration that can process 8 threads per core (or 96 threads simultaneously per clock cycle). By comparison, Intel's E5-2697) has 12 cores but can only process two threads per core for a total of only 24 threads per clock cycle. This one design difference – processor efficiency – is extremely important when comparing POWER architecture to Xeon x86 architecture. A single 12 core POWER8 processor can process over three times as much data per clock cycle as compared to a 12 core Xeon E5-2697. *What this means is that it could take up to three Xeon servers to do the work of a single POWER8-based server. It also means x86 buyers would potentially need to purchase up to three times the number of software licenses when opting for a Xeon-based server solution*.

Figure 1 presents a side-by-side comparison of Intel's Xeon E5-2697 v2 versus IBM's POWER8 architecture. Especially important to note are:

1. The # of threads/core (POWER8 can process 8 threads per core per clock cycle to Intel's two threads) – this gives IBM a huge processor efficiency advantage over Xeon;
2. The amount of data that can be placed in cache (POWER8 offers over three times as much cache). It should also be noted that POWER8 can also make use of 128 MB eDRAM L4 cache that resides just off the chip. All of this close-proximity cache gives IBM's POWER8 a huge data processing speed advantage over Intel's Xeon E5 architecture; and,
3. The memory bandwidth speed (POWER8 is almost four times faster than Xeon).

Figure 1 – Key Xeon/POWER Microprocessor/Subsystem Differentiators

	Intel Xeon E5-2697v2	POWER8
Processor Speed	2.697 GHz	4.15 GHz
Cores (Single Socket)	12	12
# of Threads/Core	2	8
Max Main Memory	768 GB	1TB
Memory Controllers	1	2
On-Cip Cache		
Level 1	32 KB I + 32 KB D/core	64 KB/core
Level 2	256KB/core	512KB/core
Level 3	30 MB/chip	96MB/chip
Memory Bandwidth	59.7 GB/s	230 GB/s
Input/Output	PCIe Gen3	PCIe Gen 3
Device Accelerator	QPI	CAPI

*Source: Clabby Analytics – June, 2014*

CAPI – Helping to Create a New Generation of Large Memory and Hybrid Systems
In addition to huge performance and efficiency advantages, IBM has also "CAPI-enabled" its
POWER8 processors. With POWER8, IBM has placed PCIe Gen 3 logic directly on the chip – and
has built an interface to this logic known as the coherence attached processor interface (or CAPI).
As illustrated in Figure 2, *CAPI is a customizable hardware accelerator* that enables devices, Flash
and coprocessors to talk directly and at very high speeds with POWER8 processors. Xeon offers a
similar interface known as Quick Path Interconnect (QPI) – the primary differentiator is that CAPI
is an open interface while QPI is not.

*The reason this CAPI interface is so important is because it eliminates the need for a PCIe bridge as well
as the need to launch the thousands of operating system and driver instructions (perhaps as many as
22.5K instructions) that are run every time PCIe I/O resources need to be used. Instead, the logic for
driving I/O resides on the chip where the number of instructions are reduced and the speed of interactions
between the CPU and associated hardware devices is dramatically improved.*

*As research analysts, the reason that we find IBM's CAPI interface to be especially noteworthy is that CAPI
breaks down barriers between the CPU and underlying hardware – making it possible for large banks of
Flash memory or coprocessors such as FPGA's and GPUs (someday) to interact at high speed with
POWER8 processors. Last year we started writing about an evolving new generation of servers that has
started to show up on our radar: specialized hybrid systems such as the VelociData appliance that can
analyze streams of structured and unstructured data at line speed. This appliance features multiple types
of processors (x86 CPUs, FPGAs, and GPUs) that sustain very high throughput and process data in a
highly parallel, lightning-fast fashion. And, as a result, huge volumes of Big Data can be read – and the
extract, transform and load cycle is greatly shortened. With CAPI, IBM and its ecosystem partners
(POWER architecture is now open and now several new business partners are building new, integrated
solutions as part of the OpenPOWER Foundation) are now working to build new, innovative coprocessor
systems that blend POWER8, GPUs and FPGAs into specialized appliance solutions. And IBM claims that
some early results show that, with CAPI, POWER8 is now capable of analyzing Big Data workloads up to
1,000 times faster than competing x86-basedservers with the same amount of memory and the same
number of cores.*

Figure 2 – The New CAPI Interface
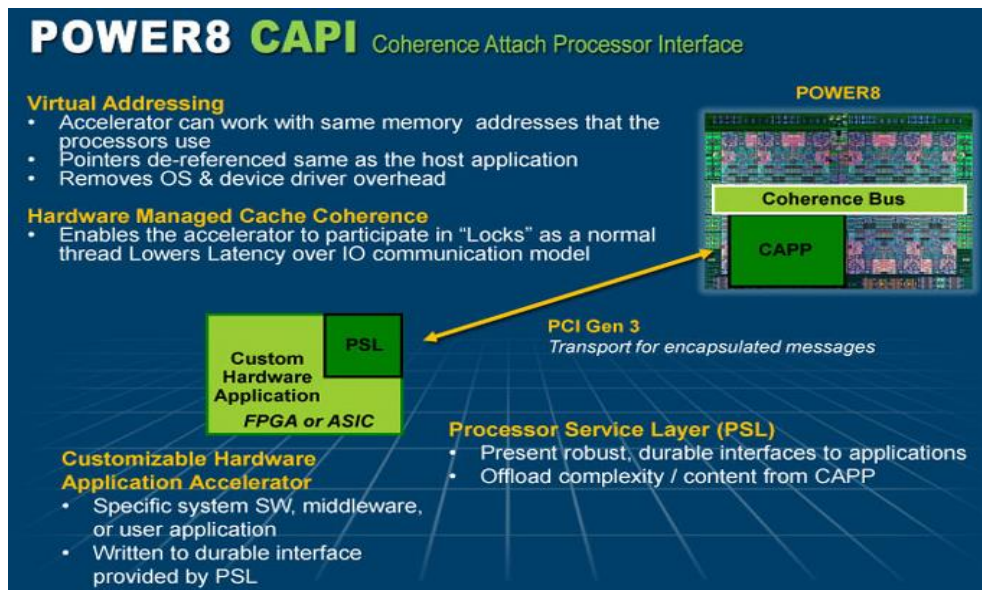
Source: IBM Corporation – June, 2014

Another Important Differentiator: Transactional Memory

Another important differentiator between POWER8 and Xeon is the availability of transactional memory on POWER8. Transactional memory is a style of execution that focuses on improving chip performance by reducing the need to handle software locking. Usually, when a thread accesses data, it acquires a lock on that data to prevent other threads from accessing that data. Although locks are a great mechanism for protecting data, they do create a lot of overhead as large numbers of threads compete for resources. One way to reduce this overhead is to use a process known as "transactional memory". What transactional memory does is it eliminates the need to lock data every time that data is accessed. Essentially, transactional memory lets the hardware decide how to proceed (essentially becoming a dynamic conflict mechanism). With transactional memory, transactions are allowed to run without locking down resources ─ and are then checked for contention issues after completion. Further, SIMD (single instruction, multiple data) instructions get a big performance boost with transactional memory because they operate on blocks of memory – yielding better memory alignment and thus better performance.

> *The reason that this approach is important is that most of the time when a transaction completes there are no contention issues ─ so removing the need for locking can greatly speed up workload processing. POWER8 supports transactional memory; Intel's Xeon v2 processors do not.*

New Machines – New Databases

Up to this juncture this *Research Report* has focused on the differences between POWER8 and Xeon v2 architectures from a processor and subsystems perspective. It should be evident to readers that with the introduction of POWER8 processors – the nature of servers is changing. By removing memory constraints and by finding new ways to link coprocessors and Flash memory to these new generation processors, new types of in-memory and hybrid configurations are starting to emerge. And, to exploit these new configurations, big changes are happening in the way that databases are being structured to deal with data analysis.

Last year *Clabby Analytics* wrote a report that discussed some of the differences between IBM's DB2 BLU, Oracle's Exadata and SAP's HANA database environments. What this report showed was that all three vendors build database environments designed to help accelerate the processing of Big Data – but that there are some very significant differences in the way that each vendor organizes/queries data – and in related system designs:

- IBM's approach to accelerating Big Data processing uses an innovative new technique known as "DB2 BLU Acceleration" to speed Big Data processing. Using a columnar approach, "BLU" quickly whittles down the size of a Big Data database to isolate relevant data, in effect speed reading large databases. This enables BLU to achieve a 10-50X performance advantage over traditional databases. The BLU approach also features database compression, the ability to read compressed data in memory – and it exploits IBM's x86 and POWER7+ and POWER8 systems designs.

- SAP's HANA relies on placing large amounts of columnar data in main memory where the whole database can be analyzed in real time. HANA also compresses data by up to 20X, but needs to decompress it to enable query processing. SAP and IBM's Power Systems organization have just entered into a testing phase with HANA – an arrangement that we believe will greatly accelerate the speed of HANA in-memory processing as well as enable many more HANA queries to be launched and processed in unison.

- Oracle's own Website describes Exadata Database Machine as "combining massive memory and low cost disks to deliver the highest performance and petabyte scalability at the lowest cost". This architecture closely resembles an Oracle real application cluster (RAC) that has been packaged as an appliance with storage that uses the Oracle database along with in database advanced analytics. It does not exploit columnar data; does not read compressed data; and it compression facilities lag IBM's DB2.

> *Each of these databases has been designed to place massive amounts of data in memory where that data can be read extremely quickly. IBM's DB2 BLU is all about building a database compressed in memory that can be exploited by POWER or x86 processors. SAP's HANA and Oracle's Exadata solutions are x86-only solutions. All of these solutions can help organizations analyze Big Data quickly – the questions that need to be answered when selecting one of these databases, however, are: 1) does this database optimally exploit underlying resources (or does the enterprise need to pay for additional resources to get the job done)? and, 2) can the database of choice meet query processing requirements in the time required? Given the POWER8 advantages in processor efficiency and memory/I/O subsystem speed, Power Systems should be strongly considered when running these kinds of data-intensive applications.*

New "Solutions Sets" to Simplify Database/Analytics Deployment

To simplify the deployment of database and analytics products on Power Systems, IBM recently announced the availability of three specially tuned and optimized "solution sets" for Power Systems. These include the:

1) IBM Solution for BLU Acceleration;
2) IBM Solution for Analytics; and,
3) IBM Solution for Hadoop.

The IBM Solution for BLU Acceleration helps speed the deployment of Big Data databases on Power Systems. The IBM Solution for Analytics accelerates the "speed of insight" for data driven analytical, computational and cognitive workloads through integration (by blending Cognos, SPSS and DB2 with BLU Acceleration). And the IBM Solution for Hadoop speed Hadoop set-up and deployment on Power Systems.

Summary Observations

When selecting computer systems, the primary goal of information technology (IT) decision makers should be to pick the computer system best suited to most efficiently execute assigned workloads. By choosing the right information systems IT executives can lower computing costs (because fewer computing systems are needed, and because fewer software licenses will be required). Further, more efficient systems often yield faster computing results (a Quality-of-Service [QoS] consideration). ___Accordingly, the choice of infrastructure (microprocessors, system designs, systems software, etc.) matters tremendously___.

The major differentiators when comparing Xeon 35 v2 processors with POWER8 microprocessors can be found in performance and efficiency:

- *Performance* – POWER8's clock speed is almost 25% faster than the E5-2692; it has access to three times more on-chip cache – and data in memory can be fed to POWER8 four times faster than its Xeon competitor. This kind of optimization makes POWER8 a formidable competitor – especially when running data-intensive applications.
- *Efficiency* – When running the same workload on a POWER8 as compared with a Xeon E5 competitor, expect more work to be processed per clock cycle (to be precise, expect three

times as much work to be processed per clock cycle). Because POWER8 can process more work more quickly, expect to have to use fewer POWER8-based systems to handle an identical workload (or expressed differently, expect to need to purchase up to three Xeon E5-based servers to handle the same amount of work as a POWER8-based server. Also expect to spend up-to three times more money for additional software licenses).

IBM's POWER8 announcement focused quite a bit on the performance and efficiency benefits that can be derived by adopting POWER8-based scale-out Power Systems. But we are also intrigued by some of the new innovations taking place within the Power Systems. At *Clabby Analytics*, we believe that, over the next several years, hundreds of new hybrid coprocessor products will come to market – bringing new innovations and new performance deltas along with them. We've already described three of these new hybrid architectures in previous *Clabby Analytics* reports on IBM's The Now Factory, the VelociData appliance and IBM's DB2 Analytics Accelerator environment. IBM's new CAPI interface to POWER8 will serve as a means to help third party business partners integrate new types of hardware (such as FPGAs and someday GPUs) with POWER8-based servers. Stay tuned for a whole new generation of really fast data appliances that will make it fast and simple to process vast volumes of Big Data.

In the end, both POWER8 and Xeon v2 x86-based architectures perform well in serial, parallel and data-intensive environments. But enterprises that need results faster – and enterprises that are interested in reducing their computing costs – should more closely evaluate POWER8-based servers. These servers can deliver results more quickly than Intel E5 x86-based solutions – which can lead to better service level performance as well as potential competitive advantages. And these servers can reduce computing costs by processing work more efficiently while requiring fewer servers and licenses to do so.

> *When all is said and done, the choice of which infrastructure to use matters. It matters a lot…*