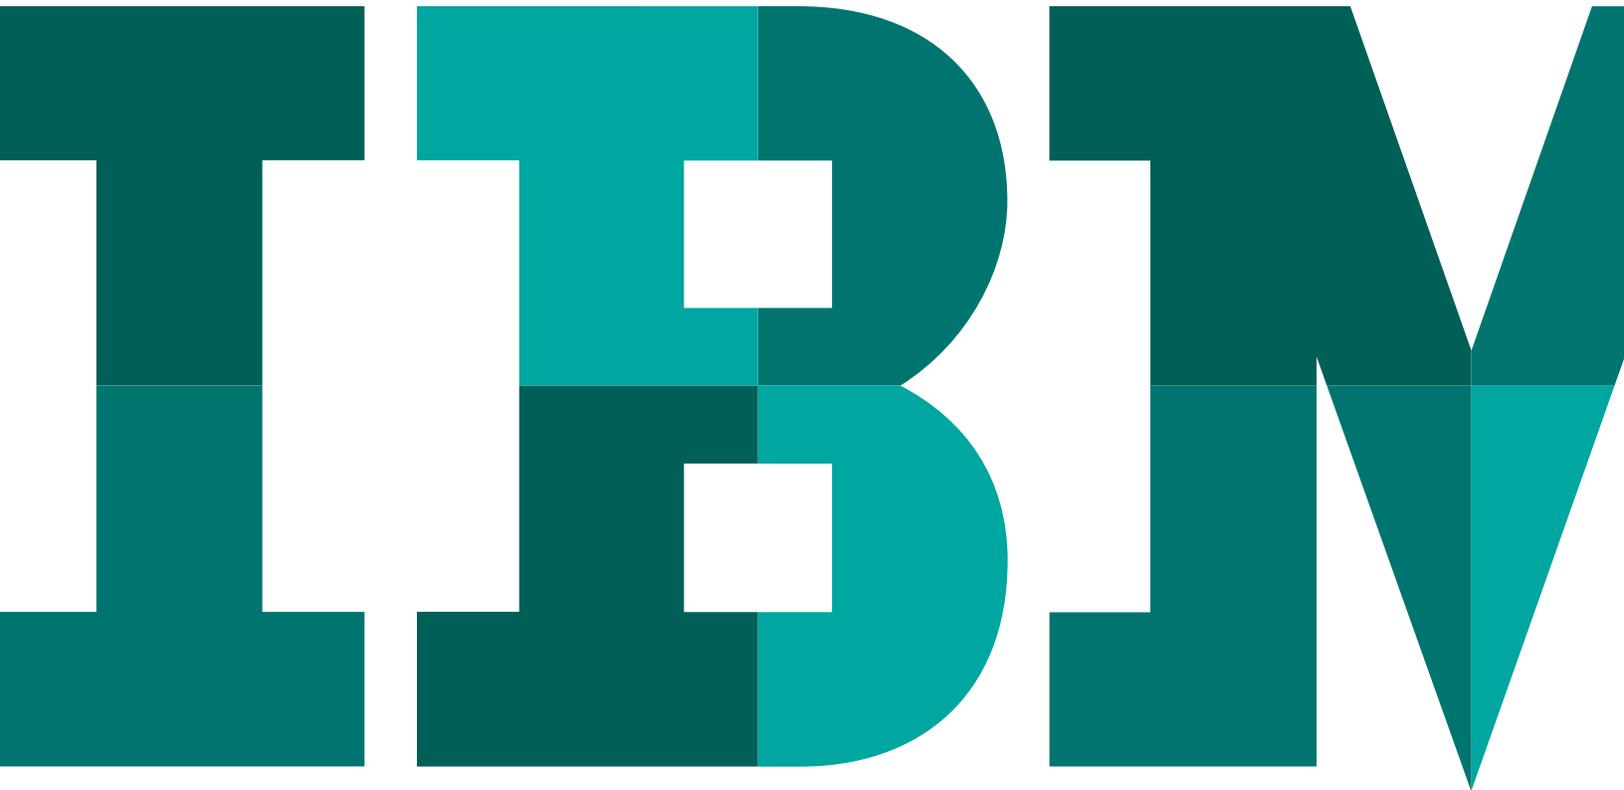


IBM POWER8 processor and memory buffer products

*Leadership performance and open innovation for
big data and cloud*



The vision today for the IBM POWER8 processor

Each day, we create 2.5 quintillion bytes of data that is generated by various sources — from climate information to posts on social media sites, and from purchase transaction records to healthcare medical images. Today’s world revolves around data, and it is “in motion,” ever evolving, ever changing, and ever growing.

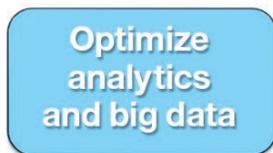
The ability to access, use, and analyze this data in order to ultimately solve problems and to reach solutions is a universal challenge. This mass of ever-growing data is typically unstructured. The process of scanning, organizing and analyzing large unstructured data sets requires high-throughput compute resources that contain high

bandwidth access to large memory capacities, and strong thread performance to address serial code segments.



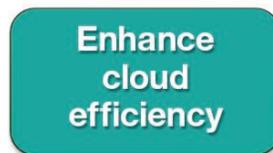
Delivering significant performance gains is increasingly difficult. While technology continues to scale to smaller dimensions, achieving performance gains requires higher levels of innovation. By using a scalable design that is combined with extensive and open enablement, IBM offers answers to the problem of big data.

IBM® POWER8™ processor and memory buffer products deliver the first truly integrated designs to address the most compelling challenges, as highlighted in Figure 1.



Leadership performance

- Increased core throughput of single thread, simultaneous multithreading levels (SMT2, SMT4 and SMT8)
- A large increase in per-socket performance
- More robust multsocket scaling



System innovation

- A higher capacity cache hierarchy and a highly threaded processor
- Enhanced memory bandwidth, capacity and expansion
- Dynamic code optimization
- Hardware-accelerated virtual memory management



Open innovation on POWER8, by offering:*

- POWER8 processor and memory buffer modules
- Open system software
- Coherent Accelerator Processor Interface (CAPI)
- Microarchitecture for your IBM POWER® chip design

* Some items require a license, purchase or both

Figure 1: POWER8 vision

Nuts and bolts: The capabilities of the POWER8 processor

The POWER8 processor (Figure 2) uses IBM's 22nm silicon-on-insulator (SOI) technology, with industry-leading speed, density and reliability. The technology density enables over 4.2 billion transistors. The design includes 15 wiring layers to create high-bandwidth, low-latency data superhighways, making it uniquely optimized for big data server applications. With embedded dynamic random access memory (eDRAM) for high-density, low-power on-chip data storage, the processor enables greater than 2.5 times the cache of a typical high-end x86 server chip.

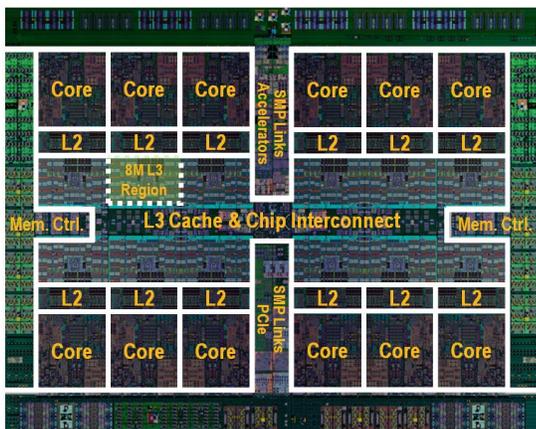


Figure 2: POWER8 processor

The POWER8 processor has several features to help data centers manage their data requirements at incredible speed.

This processor has greater capacity with 12 cores, which is 50 percent more than the IBM POWER7® processor. The cores are stronger with eight threads each, which is four times the number of data processing threads than an x86 core has. These cores have 16 execution pipelines for massive data crunching. With more than 100 MB of caches, the POWER8 processor supports the large data sets that are required by today's big data and analytics workloads. The IBM-released package for the POWER8 processor supports up to 115 GBps, for more than two times the ability of the current x86 technology to feed data into its cores.

The open interfaces of the POWER8 processor include up to 96 GBps (64 GBps in the released package) of integrated Peripheral Component Interconnect Express (PCIe) Gen3 and CAPI. The CAPI capabilities support coherent attachment to the processor, enabling low latency flash memory, tight network integration, and Field Programmable Gate Array (FPGA) acceleration.

In addition, the POWER8 processor supports flat two-hop symmetric multiprocessing (SMP) interconnect. For enhanced energy efficiency, the processor also consists of an on-chip microcontroller for power management control.

In total, POWER8 delivers three times the socket performance of its predecessor POWER7.

POWER8 Core Capabilities

Although most multi-core processors focus on core count increase at roughly constant core performance, the POWER8 core (Figure 3) has an enhanced microarchitecture that doubles the POWER7 per-core throughput.



Figure 3: POWER8 core

The POWER8 core has execution improvements, larger caching structures, wider load and storage capacities, enhanced prefetch and greater performance when compared with the POWER7 core. Starting with the execution improvements, the POWER8 core has a higher SMT level (SMT8 from SMT4), eight instructions dispatched per cycle, ten operations issued to execution pipelines and larger (four 16-entry) issue queues. It includes 16 execution pipes, consisting of: two fixed-point units (FXUs), two load/store units (LSUs), two load units (LUs), four floating-point units (FPUs), two vector (VMX) units, one Crypto, one decimal floating-point unit (DFU), one conditional register unit (CRU) and one branch execution unit (BRU). This core supports larger global completion and load or store instruction reordering. Plus, it has improved branch prediction and unaligned storage access.

The POWER8 core has twice the L1 data cache (64 KB) and supports twice the outstanding data cache misses and four times the translation cache over the POWER7 core. The POWER8 core includes a 64B (previously 32B) L2 to L1 reload data bus and twice the data cache to execution data flow. The enhanced prefetch of the POWER8 core includes instruction speculation awareness, data prefetch depth awareness, adaptive bandwidth awareness and topology awareness. Finally, compared to the POWER7 core, the POWER8 core has greater performance at all levels, with almost 1.6 times the single thread performance and twice the core throughput.

Memory and cache optimizations

The IBM POWER8 cache hierarchy was designed for big data workloads and doubles the cache capacity and bandwidth of its predecessor, the POWER7 processor. The POWER8 cache hierarchy introduces up to an extra 64 MB of shared L4 cache per processor by using the attached new memory buffer chips. The POWER8 cache hierarchy is aligned end to end for high data bandwidth. It includes double-wide data flows that extend through the L1, L2 and L3 caches and the on-chip interconnect to the memory read/write interfaces. See Table 1.

	POWER7	POWER8
L1 data capacity/core	32 kB	64 kB
L2 capacity/core	256 kB	512 kB
L3 capacity/core	4 MB	8 MB
L4 capacity/chip	N/A	64 MB
L2 bandwidth/core at 4 GHz	256 GBps	512 GBps
L3 bandwidth/core at 4 GHz	128 GBps	256 GBps
Total L3 interconnect bandwidth	512 GBps	1382 GBps

Table 1: POWER8 cache attributes

In addition to the cache enhancements, the POWER8 solution provides a balanced system in support of the increased computing capacity by improving the I/O subsystem and the capabilities of the memory buffer chips. The end-to-end data and coherence bandwidth of the system is more than twice the bandwidth of the POWER7 system.

Memory organization and memory buffer chip

In the released package, the POWER8 processor memory organization provides the potential for up to four high-speed channels that connect to the new memory buffer chips. These channels each run up to 9.6 GBps for up to 115 GBps of bandwidth to memory. Each memory buffer chip then has four DDR ports with the potential to yield 205 GBps at the DRAM and up to a total of 1024 GB memory capacity per fully configured processor socket. The new memory buffer chip and specific design features are essential elements of the performance, capacity, and scalability improvements in the POWER8 processor. Figure 4 shows the base building blocks of the memory architecture.

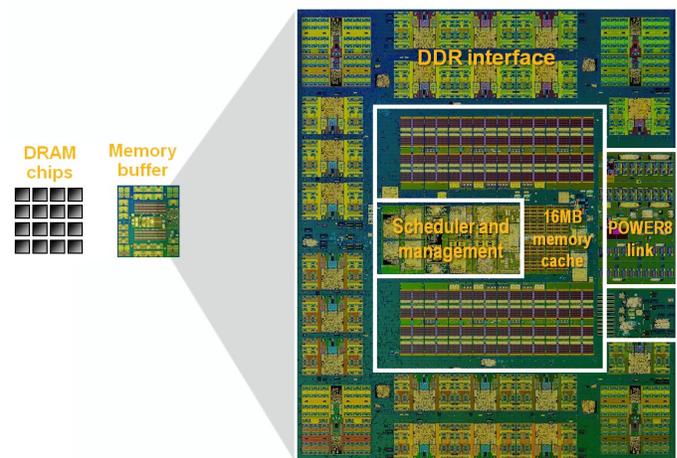


Figure 4: POWER8 processor memory organization

The memory architecture in the POWER8 processor has the following design features and benefits:

- Intelligence moved into memory
 - Scheduling logic, caching structures
 - Energy management and reliability, availability and serviceability (RAS) decision point
- Processor interface
 - A 9.6 GBps high-speed interface
 - More robust RAS
 - Dynamic lane isolation and repair
 - Extensible for innovation build-out
- Performance value
 - About a 20 percent latency reduction
 - Additional energy and latency benefits from cache
 - Improved scheduling and memory usage
 - A 90 percent processor-to-memory channel efficiency

IO subsystem and SMP interconnect

The POWER8 processor integrates PCIe Gen3 to increase overall I/O bandwidth and to reduce latency. Integrating PCI also reduces system power consumption and frees up card space. The 32 lanes of integrated PCIe Gen3 with three PCI host bridges provide 64 GBps of I/O bandwidth. They achieve a direct memory access (DMA) read latency of less than one-third of what the POWER7 processor could achieve with a discrete I/O hub chip.

The on-node and off-node symmetric multiprocessing (SMP) buses enable 494 GBps of chip-to-chip communication that includes data, command, control, error correction and sparing.

Power management—It all starts with the silicon

POWER8 also establishes a new baseline for the artful balance between high performance and energy efficiency — a critical factor for stand-alone, Internet-scale, and large-scale distributed computing systems. POWER8 features a new on-chip power management microcontroller which monitors a variety of processor and system variables, and then responds by dynamically setting the frequency and voltage levels of each individual core to match workload demand. With a variety of power management modes to select from, POWER8 systems provide the flexibility to meet growing demands of datacenter energy constraints at the individual application and server level.

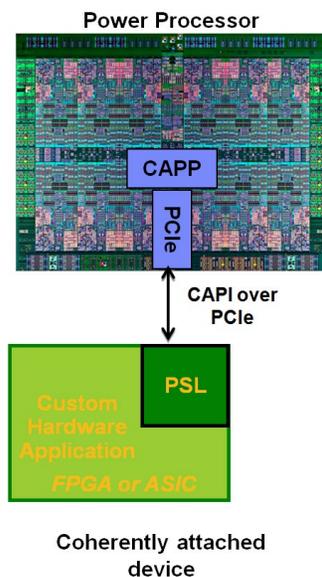
Open ecosystem enablement and acceleration

Because server workloads will continue to evolve, POWER8 established an open standard that enabled ongoing off-processor innovation and adaptation in POWER8 systems from various industry sources. In addition, many workloads can benefit from encapsulating specific algorithms in hardware to support the general-purpose cores for a heterogeneous computing solution.

 To meet these needs, the POWER8 system introduces the Coherent Accelerator Processor Interface. CAPI provides the capability for off-chip accelerators to participate in the system-memory coherence protocol as a peer of other caches in the system. CAPI also allows these accelerators to use effective addresses to reference data structures just as an application that is running on the cores does. These accelerators can be plugged into PCIe slots and implemented in FPGAs or application-specific integrated circuit (ASIC) chips.

The CAPI architecture (see Figure 5) is implemented in three pieces. First, the Coherent Attached Processor Proxy (CAPP) unit extends coherency to an attached device. An on-processor directory responds on behalf of an off-chip device.

Second, the coherency protocol of the CAPI is tunneled over standard PCIe. Using industry-standard PCI physicals eliminates the need for special I/Os and protocol logic because CAPI uses standard posted write and non-posted reads. It also reduces the complexity and bandwidth requirements of the attached device. Third, the processor service layer (PSL) provides the translation and interrupt services for the application.



CAPI enables an attached device to be a peer to the processor. In doing so, it simplifies the programming model between applications and enables the device to use the same effective address as an application that is running in the processor. It also eliminates cumbersome I/O device driver requirements. (Pinned memory is not required.)

Finally, CAPI enables a customizable hardware application accelerator for specific system software, middleware or user application. It is written to a durable interface that is provided by the PSL.

Summary

IBM has implemented significant technical advances in the IBM POWER8 processor. It includes 22nm technology, core execution improvements, memory organization, power management, and open I/O that are available in the combined processor and memory buffer chips. These features deliver a design that is optimized for big-data, analytics, and cognitive workloads, while supporting efficient cloud operating environments. In addition, the newly introduced CAPI framework and open software stack components offer off-chip innovation and acceleration opportunities. The open availability to now access, license and use these POWER8 chip sets provides game-changing technology and IP to the enterprise system community.

Figure 5: Coherent Accelerator Processor Interface

For more information

For more information about IBM POWER8 module product offerings, see the following documents by means of the IBM webpage link below (sign-on required):

ibm.com/technologyconnect/index

- POWER8 Processor data sheet for the Single-Chip Module, Revision Level DD 2.X
- POWER8 Memory Buffer data sheet

To learn more about IBM POWER8 chip sets, contact your IBM sales representative or send email to openpower@us.ibm.com.



© Copyright IBM Corporation 2014

IBM Corporation
IBM Systems and Technology Group
2070 Route 52, Bldg. 330
Hopewell Junction, NY 12533-6351

Produced in the United States of America
July 2014

IBM, the IBM logo, ibm.com, POWER, POWER 7, POWER7+, and POWER8 are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.



Please Recycle
