Maximizing the Value of an IBM POWER7 and IBM POWER7+ Environment through Tuning and Optimization

IBM Redbooks Solution Guide

Strategies for optimizing and tuning application code to run on IBM POWER7® and IBM POWER7+™ processor-based systems can be invaluable to your environment and to your business. They can substantially improve the performance of the applications that run on these systems. Optimizing and tuning your IBM Power Systems™ environment can be an important step in meeting your critical business needs. Optimized systems will deliver the performance to meet your current requirements and your future growth needs. By using the strategies provided in this solution guide, you can maximize the return on your hardware investment with minimal effort. These strategies can provide an avenue to deliver continuing, long-term value over the life of your system.

The information in this solution guide is drawn from application optimization efforts across many types of code running on the IBM AIX® and Linux® operating systems. It focuses on the more pervasive performance opportunities that are identified and how to capitalize on them. This technical information was developed by IBM domain experts and is directed to IBM presales organizations in support of Power System products, such as the IBM Power 780 (Figure 1).

Figure 1. IBM Power 780 server

Did you know?

Trends in processor design are making it more important than ever to consider improving application performance. The focus of processor design has shifted to delivering multiple cores per processor chip and to delivering more hardware threads in each core (known as *simultaneous multithreading* (SMT) in IBM Power Architecture® terminology). Some of the best opportunities for improving application performance are in delivering scalable code by having an application effectively use multiple concurrent threads of execution. Another trend is support for larger page sizes. The IBM Power Architecture provides support for multiple virtual memory page sizes, which provides performance benefits to an application because of hardware efficiencies that are associated with larger page sizes.

Business value

You can follow simple strategies and techniques to optimize your POWER7 environment and to analyze and maximize system performance. These strategies and techniques can be invaluable and offer the following advantages:

- Substantially improve the performance of the application that is being optimized for POWER7
- Typically carry over improvements to systems that are based on related processor chips
- Improve performance on other platforms

Optimization guidelines are provided in the following categories:

- *Lightweight tuning and optimization guidelines*, which include simple, prescriptive steps for tuning application performance on POWER7. Most can be carried out without modifying application source code.

- *Deployment guidelines*, which include steps for configuring POWER7 to optimize performance by making choices among the deployment alternatives.

- *Deep performance optimization guidelines*, which include tools and strategies for identifying and fixing application bottlenecks. This analysis requires more familiarity with performance tools and analysis techniques.

These guidelines can be applied to all IBM POWER® generations, including the newest IBM POWER7+ processor. The concise introductory guidelines of this solution guide and the comprehensive nature of *POWER7 and POWER7+ Optimization and Tuning Guide*, SG24-8079, make these valuable resources in your IBM Power Systems environment.

Solution overview

The techniques to optimize your POWER7 environment and to analyze and maximize system performance capitalize on the capabilities and features of the following products:

- The IBM POWER7 processor
- The IBM POWER Hypervisor™
- IBM AIX, including Active System Optimizer (ASO), Dynamic System Optimizer (DSO), and AIX memory allocation (malloc)
- Linux, which is optimized for Power Architecture

The IBM POWER7 processor

Several capabilities and features of the POWER7 processor are key to system optimization. POWER7 offers the following most important, yet simple features for performance tuning:

- Multiple page size support feature

  Power Architecture supports multiple virtual memory page sizes, which in turn, provide performance benefits to an application because of hardware efficiencies that are associated with larger page sizes. Large pages provide several technical advantages such as the following examples:

  o Reduced page faults and Translation Lookaside Buffer (TLB) misses

    A single large page that is being constantly referenced remains in memory, eliminating the possibility of swapping out several small pages.

  o Unhindered data prefetching

    A large page enables unhindered data prefetch, which is constrained by page boundaries.

  o Increased TLB Reach

    This feature saves space in the TLB by holding one translation entry instead of *n* entries, which increases the amount of memory that can be accessed by an application without incurring hardware translation delays.

  o Increased Effective to Real Address Translation (ERAT) Reach

    ERAT on IBM POWER is a first-level and fully associative translation cache that can go directly from effective to real address. *Effective addresses* are the addresses used by the software, and *real addresses* refer to the physical memory that is assigned to the software by the system. Both the ERAT and the TLB are involved in translating addresses. Large pages also improve the efficiency and coverage of this translation cache.

- POWER7 processor and affinity performance effects

  The IBM POWER7 and POWER7+ are the latest processor chips in the Power Systems family. The POWER7 and POWER7+ processor chips are available in configurations with four, six, or eight cores per chip, as compared to the IBM POWER5® and IBM POWER6® processor chips, which have two cores per chip. Along with the increased number of cores, the POWER7 and POWER7+ processor chips implement SMT4 mode, which supports four hardware threads per core. The POWER5 and POWER6 support only two hardware threads per core. Each POWER7 and POWER7+ processor core supports running in single-thread mode with one hardware thread, in SMT2 mode with two hardware threads, or in SMT4 mode with four hardware threads.

  Each SMT hardware thread is represented as a logical processor in AIX or Linux. When the operating system runs in SMT4 mode, it has four logical processors for each dedicated POWER7 and POWER7+ processor core that is assigned to the partition. To gain full benefit from the throughput improvement of SMT, applications must use all of the SMT threads of the processor cores.

  Each POWER7 and POWER7+ chip has memory controllers that allow direct access to a portion of the memory dual inline memory module (DIMMs) in the system. Any processor core on any chip in the system can access the memory of the entire system. However, it takes longer for an application thread to access the memory that is attached to a remote chip than to access data in the local memory DIMMs.

  Affinity effects are related to the efficient use of the caches on a POWER7 and POWER7+ chip and to the memory that is local to each chip. Software threads that access the same data are best run together on the SMT4 threads of a single core and on the cores of a single chip. All of the data that is accessed from a chip should be in local memory and not in remote memory. For an example of the use of SMT4 mode, see the usage scenario in this solution guide.

The IBM POWER Hypervisor

The IBM POWER Hypervisor manages the virtualization of processor cores and memory for the operating system. It also ensures that the affinity between the processor cores and memory that a logical partition (LPAR) is using is maintained as much as possible. However, application designers must also consider affinity issues. Another key aspect of POWER Hypervisor is the impact of application thread and data placement on the cores and the memory that is assigned to the LPAR that the application is running in.

IBM PowerVM® Hypervisor and the AIX operating system (version AIX V6.1 TL 5 and later) on POWER7 implement enhanced affinity in several areas. This feature achieves optimized performance for workloads that are running in a virtualized shared processor LPAR (SPLPAR) environment. These areas can include virtual processors, LPAR page table sizes, and placing LPAR resources to attain higher memory affinity.

AIX: Active System Optimizer, Dynamic System Optimizer, and AIX malloc

AIX benefits from the following optimization and tuning techniques:

- AIX Active System Optimizer (ASO), the Dynamic System Optimizer (DSO)

  Workloads are becoming increasingly complex. Typically, they involve a mix of single-thread and multithread applications with complex interactions that vary over time. The servers that host these workloads are continuously evolving to support an ever-increasing demand for processing capacity and flexibility. Optimizing such an environment often requires excessive amounts of time and highly specialized skills. Further, manual tuning is static in nature, and systems must be retuned on occasion. ASO and DSO help to optimize the operating system and server autonomously.

  ASO provides two optimization strategies:

  o   Cache affinity optimization
  o   Memory affinity optimization

  DSO (built on the ASO framework) adds two more optimization strategies to the ASO framework:

  o   Large page optimization
  o   Memory prefetch optimization

The ASO framework (Figure 2) continuously monitors and analyzes how current workloads impact the system. It then uses this information to dynamically configure the system to optimize for current workload requirements. The ASO framework is transparent. The administrator is not required to continuously monitor its operations. ASO uses information from the AIX kernel and the POWER7 performance monitoring unit (PMU) to perform long-term runtime analysis to improve workload performance.
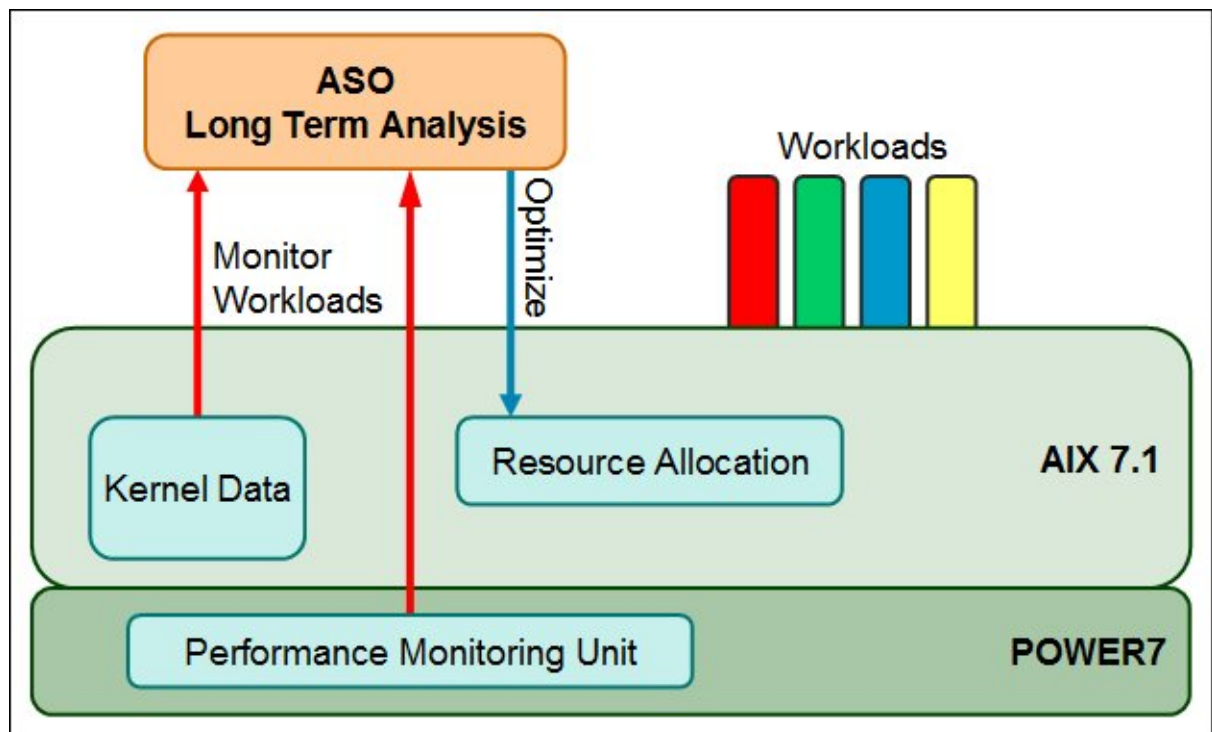

Figure 2. Basic ASO architecture that shows optimization flow on a POWER7 system

The primary design goal of ASO/DSO is to act only when it is reasonably certain that the result is an improvement in workload performance.

● AIX memory allocation (malloc)

AIX malloc is another optimization and tuning technique for AIX. The AIX operating system offers various memory allocation packages (the standard `malloc()` and related routines in the C library). The default package offers good space efficiency and performance for single-thread applications, but it is not a good choice for the scalability of multithread applications. Choosing the correct malloc package on AIX is important for performance. Even Java applications can extensively use malloc through Java Native Interface (JNI) code or internally in the Java runtime environment (JRE).

Fortunately, AIX offers several different memory allocation packages that are appropriate for different scenarios. These packages are chosen by setting environment variables and do not require any code modification or rebuilding of an application. Choosing the best malloc package requires an understanding of how an application uses the memory allocation routines. To learn how to easily collect the required information, see Appendix A, "Analyzing malloc usage under AIX" in *POWER7 and POWER7+ Optimization and Tuning Guide*, SG24-8079. After the data collection, experiment with the various alternatives, alone or in combination.

The following packages are some alternatives that deliver high performance:

o   **The pool malloc option**: The pool front end to the malloc subsystem optimizes the allocation of memory blocks of 512 bytes or less. It is common for applications to allocate many small blocks, and pools are particularly space-efficient and time-efficient for the allocation pattern. Thread-specific pools are used for multithread applications. The pool malloc is a good choice for both single-thread and multithread applications.

o   **The multiheap malloc option**: The multiheap malloc package uses up to 32 separate heaps, reducing contention when multiple threads attempt to allocate memory. It is a good choice for multithread applications.

Using the pool front end malloc and the multiheap malloc in combination is a good alternative for multithread applications. Small memory block allocations, which are typically the most common type, are handled with high efficiency by the pool front end. Larger allocations are handled with good scalability by the multiheap malloc. A simple example of specifying the pool and multiheap combination is by using the environment variable setting:

```
MALLOCOPTIONS=pool,multiheap
```
For more information about using AIX malloc, see the usage scenarios in this solution guide.

Linux: Optimized for Power Architecture

A solid choice for running enterprise-level workloads on POWER7 is Linux. Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server (SLES) are optimized and targeted for the Power Architecture. These operating systems take full advantage of the specialized features of Power Systems. RHEL6 GA and SLES11 SP1 are the minimum supported versions to fully use POWER7 technologies and systems.

Both RHEL and SLES provide the tools, kernel support, optimized compilers, and tuned libraries for IBM POWER7 Systems™. The Linux distributions provide excellent performance, and more application and customer-specific tuning approaches are available. IBM provides several packages, tools, and extensions that provide for more tuning, optimization, and products for the best possible performance on POWER7. The typical Linux open source performance tools that Linux users are comfortable with are available on IBM PowerLinux™ systems.

Solution architecture

This section describes the architecture of the POWER7 processor and its capabilities for multi-core and multithread scalability.

Architecture of the POWER 7 processor

The POWER7 processor is manufactured with IBM 45 nm Silicon-On-Insulator (SOI) technology. Each chip is 567 mm^2 and contains 1.2 billion transistors. The POWER7 processor chip (Figure 3) contains eight cores. Each core has its own 256 KB L2 and 4 MB L3 (embedded dynamic random access memory (DRAM)) cache, two memory controllers, and an interconnection system that connects all components within the chip. The interconnect also extends through module and board technology to other POWER7 processors, DDR3 memory, and various I/O devices. The number of memory controllers and cores that are available for use depends on the POWER7 system.
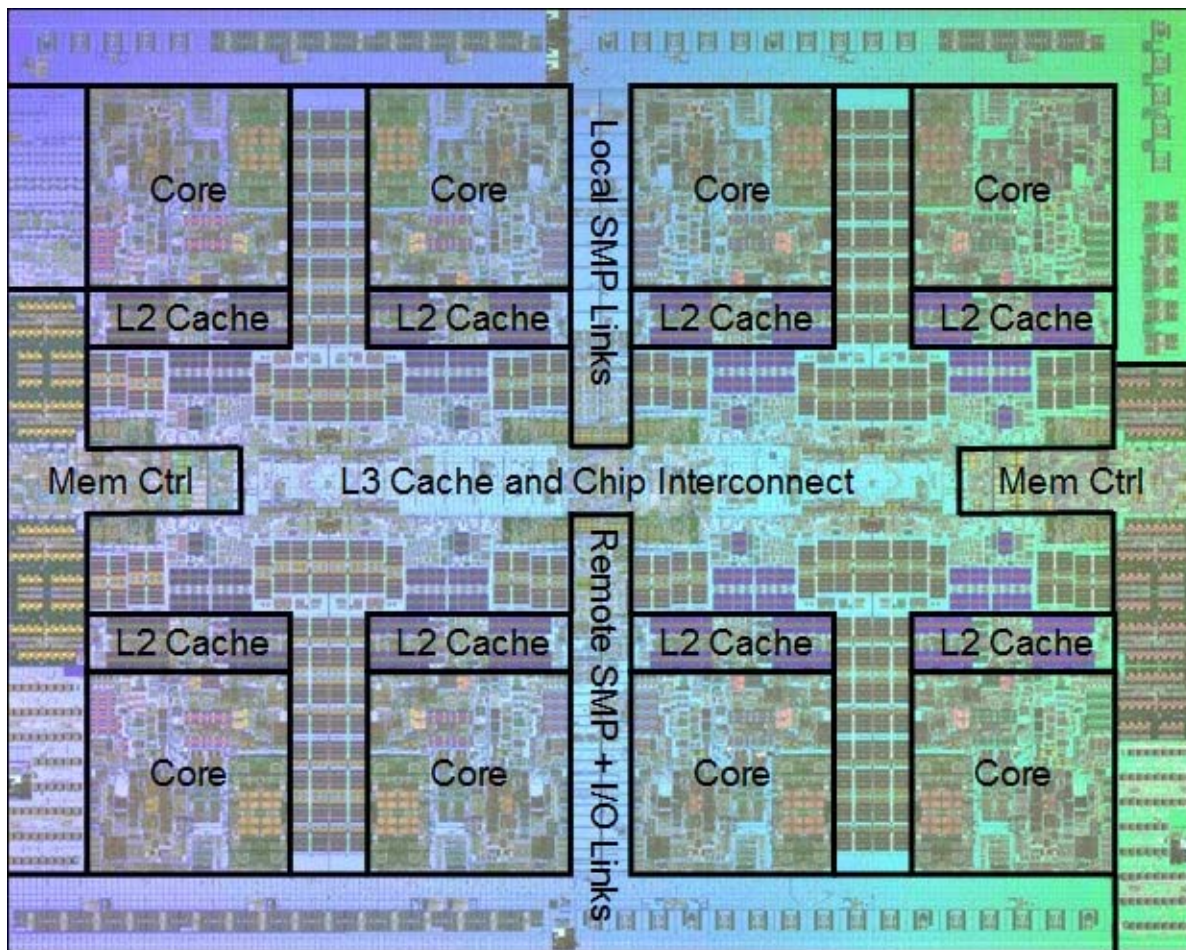

Figure 3. The POWER7 processor chip

Each core is a 64-bit implementation of the IBM Power ISA (Version 2.06 Revision B) and has the following features:

- Multithread design that supports up to a four-way SMT
- 32 KB, four-way set-associative L1 i-cache
- 32 KB, eight-way set-associative L1 d-cache
- 64-entry ERAT for effective-to-real address translation for instructions (2-way set associative)
- 64-entry ERAT for effective-to-real address translation for data (fully associative)
- Aggressive branch prediction that uses local and global prediction tables with a selector table to choose the best predictor
- 15-entry link stack
- 128-entry count cache
- 128-entry branch target address cache
- Aggressive out-of-order execution
- Two symmetric fixed-point execution units
- Two symmetric load/store units, which can also run simple fixed-point instructions
- An integrated, multipipeline vector-scalar floating point unit that supports up to eight flops per cycle and that runs the following Scalar and Single Instruction Multiple Data (SIMD)-type instructions:
    - o The Vector Multimedia Extension (VMX) instruction set
    - o The Vector Scalar Extension (VSX) instruction set
- Hardware data prefetching with 12 independent data streams and software control
- Hardware decimal floating point (DFP) capability
- Adaptive power management

The POWER7 processor is designed for system offerings from 16-core blades to 256-core drawers. It incorporates a dual-scope broadcast coherence protocol over local and global symmetric multiprocessor (SMP) links to provide superior scaling attributes.

The POWER7+ processor is the same POWER7 processor core with new technology, including more on-chip accelerators and an extra L3 cache. No new instructions are in POWER7+ over POWER7. POWER7+ differs from the POWER7 processor in that it is manufactured with the following features:

- 32-nm technology
- A 10 MB L3 cache per core
- On-chip encryption accelerators
- On-chip compression accelerators
- On-chip random number generators

Usage scenarios

This section includes examples of optimization and tuning guidance. For more examples, see *POWER7 and POWER7+ Optimization and Tuning Guide*, SG24-8079.

Usage scenario 1: Memory allocator suboptions

The following use cases relate to memory allocation and can be used to set up your environment:

- For a 32-bit, single-thread application, use the default allocator.

- For a 64-bit application, use the Watson allocator.

- Multithread applications use the multiheap malloc option. Set the number of heaps proportional to the number of threads in the application.

- For single-thread or multithread applications that make frequent allocation and deallocation of memory blocks smaller than 513, use the pool malloc option.

- For a memory usage pattern of the application that shows high usage of memory blocks of the same size (or sizes that can fall to common block sizes in the buckets option) and sizes greater than 512 bytes, use the malloc buckets option.

- For older applications that require high performance and do not have memory fragmentation issues, use malloc 3.1.

- Ideally, the Watson allocator, with the multiheap malloc and pool malloc options, are good for most multithread applications. The pool front end is fast and scalable for small allocations. The multiheap malloc option ensures scalability for larger and less frequent allocations.

- If you notice high memory usage in the application process even after you run `free()`, try using the disclaim option.

For more information, see *POWER7 and POWER7+ Optimization and Tuning Guide*, SG24-8079.

Usage scenario 2: Tuning to capitalize on hardware performance features

For almost all applications, using 64-KB pages is beneficial for performance. Newer Linux releases (RHEL5, SLES11, and RHEL6) default to 64-KB pages, and AIX defaults to 4-KB pages. Applications on AIX enable 64-KB pages through one, or a combination, of the following methods:

- Using an environment variable setting:
  ```
  LDR_CNTRL=TEXTPSIZE=64K@DATAPSIZE=64K@STACKPSIZE=64K@SHMPSIZE=64K
  ```

- Modifying the executable file as follows:
  ```
  ldedit –btextpsize=64k –bdatapsize=64k –bstackpsize=64k <executable>
  ```

- Using linker options at build time:
  ```
  cc –btextpsize:64k –bdatapsize:64k –bstackpsize:64k ...
  ld –btextpsize:64k –bdatapsize:64k –bstackpsize:64k ...
  ```

These mechanisms for enabling 64-KB pages can be used safely when you run them on older hardware or operating system levels that do not support 64-KB pages. When the necessary support is not in place, the system defaults to using 4-KB pages.

Recent Java releases default to using 64-KB pages. For Java, the Java heap space uses 64-KB pages, which are enabled by the `-Xlp64k` option in older releases (a minimum Linux level of RHEL5, SLES11, or RHEL6 is required).

Larger 16-MB pages are also supported on the Power Architecture and might provide an extra performance boost when compared to 64-KB pages. However, usage of 16-MB pages normally requires explicit configuration by the administrator of the AIX or Linux operating system. The DSO facility in AIX

autonomously uses 16-MB pages without any administrator configuration, which might be appropriate for cases where a large memory space is used by an application.

For more information, see *POWER7 and POWER7+ Optimization and Tuning Guide*, SG24-8079.

Usage scenario 3: Partition sizes and affinity with power dedicated LPARs

Consider a case in which you are running four instances of IBM WebSphere® Application Server on a partition of 16 cores on a POWER7 system that is running in SMT4 mode. For good affinity, each instance of WebSphere Application Server is bound to run on four of the cores of the system. Because each core has four SMT threads, each instance of WebSphere Application Server is bound to 16 logical processors. To ensure good memory and cache affinity on AIX:

1. Set the `AIX MEMORY_AFFINITY` environment variable. Typically it is set to the value `MCM`. This setting signals the AIX operating system to use local memory when an application thread requires physical memory to be allocated.

2. Start the four instances of WebSphere Application Server by running the following **execrset** commands in the order shown (first instance to fourth instance) to bind the execution to the specified set of logical processors:

   - `execrset -c 0-15 -m 0 -e`
   - `execrset -c 16-31 -m 0 -e`
   - `execrset -c 32-47 -m 0 -e`
   - `execrset -c 48-63 -m 0 -e`

   Keep in mind the following important items:

- For a particular number of instances and available cores, each instance of an application runs only on the cores of one POWER7 processor chip.

- Memory and logical processor binding is not done independently because doing it can negatively affect performance.

- The workload must be evenly distributed over WebSphere Application Server processes for the binding to be effective.

- An assumed mapping of logical processors to cores and chips is always established at startup. This mapping can be altered if the SMT mode of the system is changed by running the **smtctl -w now** command. Restart the system to change the SMT mode of a partition to ensure that the assumed mapping is in place.

For more information, see *POWER7 and POWER7+ Optimization and Tuning Guide*, SG24-8079.

Integration

The strategies in this solution guide apply to all POWER generations, including the POWER7+ processor.

Supported platforms

This section highlights the supported operating systems and other key prerequisites for Power Systems. For information about individual models, see the Power servers page at:
http://www.ibm.com/systems/power/hardware/index.html?&LNK=browse

Power Express servers

Power Express servers are excellent as reliable, secure distributed application servers, consolidation servers, or stand-alone servers for UNIX, IBM i, and Linux workloads. As 2U, 4U, or tower packages with

from 4 to 32 cores, Power Express servers provide outstanding performance and help to reduce infrastructure and energy costs.

Power Enterprise servers

Power Enterprise servers are for clients who require the ultimate in business resiliency, performance, and scalability. This class of system, which can run AIX, IBM i, and Linux, provides up to 256 POWER7 processor cores with up to 8 TB of memory. It includes the flexibility to turn processors and memory on and off as application workloads dictate.

PowerLinux servers

World-class POWER7 Systems are equipped with two sockets and up to 16 cores. These value-priced servers go head-to-head with x86 servers in terms of cost and in delivering greater performance, higher utilization, and superior availability.

High performance computing

High performance computing solutions with Power Systems that are configured into highly scalable AIX and Linux clusters offer extreme performance for demanding analytic and big data workloads. They can handle workloads that involve computational chemistry, petroleum reservoir modeling, weather forecasting, climate modeling, and financial services.

IBM PureFlex System

The IBM PureFlex™ System provides compute, storage, and networking resources in one environment that is efficient and easy to manage. IBM Flex System™ components provide an open environment of advanced networking, storage, and virtualization technologies with flexibility for various workloads.

Ordering information

Table 1 summarizes the ordering information. Most Power Systems models can be built to your specifications. For a customized quotation, call your IBM sales representative at 1-866-883-8901. For announcement letter and sales manual information for each offering in Table 1, see the IBM Offering Information page in the "Related information" section.

Table 1. Part numbers (feature codes) and descriptions for IBM Power Systems models (Part 1)

Power System model	Part number (feature code)	Charge unit description
IBM Power 710 Express	8231-E1C	This server is a 2U rack-mount server with one processor socket offering 4-core 3.0-GHZ, 6-core 3.7-GHZ, and 8-core 3.55-GHZ configurations.
IBM Power 720	8202-E4C	This server offers powerful 64-bit POWER7 processors that offer 4-core, 6-core, and 8-core configuration options; tower or rack-mount configuration; memory capacity increased up to 256 GB of memory with optional memory riser card, optionally augmented with IBM Active Memory™ Expansion.
IBM Power 730 Express	8231-E2C	This server is a 2U rack-mount server with two processor sockets offering 8-core 3.0-GHZ and 3.7-GHZ, 12-core 3.7-GHZ, and 16-core 3.55-GHZ configurations.

Table 1. Part numbers (feature codes) and descriptions for IBM Power Systems models (Part 2)

Power System model	Part number (feature code)	Charge unit description
IBM Power 740 Express	8205-E6C	This server is recommended when a solution requires high communications or I/O, or requires the maximum amount of memory available. PCIe Gen2 slots can transfer data at double the speed. The high data transfer rates that are offered by the PCIe Gen2 slots can allow higher I/O performance or consolidation of the I/O demands onto fewer adapters that are running at higher rates. This result is better system performance at a lower cost when I/O demands are high.
IBM Power 750 Express	8233-E8B	This server has POWER7 processors that offer 4-core to 32-core configuration options.
IBM Power 755	8236-E8C	This server is a 3.3-GHZ or 3.6-GHZ 32-core POWER7 processor-based server, providing four 64-bit, eight-core processor POWER7 modules with 4 MB of L3 cache/core and 256 KB of L2 cache/core.
IBM Power 770 POWER7	9117-MMC	This server is a modular system that might be configured with 1 - 4 processor drawers. A system that is configured with up to four of these drawers using 6-core SCM processors enables up to 48 processor cores that are running at frequencies up to 3.72 GHZ.
IBM Power 770 POWER7+	9117-MMD	This server is an SMP, rack-mounted server. This modular system uses one to four enclosures. Each contains four powerful POWER7+ processors and high-density memory DIMMs that use 4-Gb technology.
IBM Power 780	9179-MHC	This server is an SMP, rack-mounted server. This modular-built system uses 1 - 4 enclosures.
IBM Power 780	9179-MHD	This server is an SMP, rack-mounted server that uses one to four enclosures. Each enclosure contains four powerful POWER7+ processors and high-density memory DIMMs that use 4-Gb technology.
IBM Power 795	9119-FHB	This server is an SMP, rack-mounted server. Equipped with eight 32-core or 24-core processor books, the Power 795 server can be deployed in 24-core to 256-core, SMP configurations. It has up to 8 TB of buffered DDR3 memory and extensive I/O support.

Related information

For more information, see the following documents:

- *POWER7 and POWER7+ Optimization and Tuning Guide*, SG24-8079
  http://www.redbooks.ibm.com/abstracts/sg248079.html

- IBM Offering Information page (to search on announcement letters, sales manuals, or both):
  http://www.ibm.com/common/ssi/index.wss?request_locale=en

  On this page, enter any of the following names, select the information type, and then click **Search**. On the next page, narrow your search results by geography and language:

  - *Family 8231+01 IBM Power 710 and 730 Express Servers*
  - *Family 8202+02 IBM Power 720 server*
  - *Family 8205+01 IBM Power 740 Express Server*
  - *Family 8233+01 IBM Power 750 Express Server*
  - *Family 8236+01 IBM Power 755 Server*
  - *Family 9117+04 IBM Power 770 POWER7 Server (9117-MMC)*
  - *Family 9117+05 IBM Power 770 POWER7+ Server (9117-MMD)*
  - *Family 9179+02 IBM Power 780 Server (9179-MHC)*

- *Family 9179+03 IBM Power 780 Server (9179-MHD)*
- *Family 9119+04 IBM Power 795 Server*

- Enhanced I/O options for Power Systems
  http://www.ibm.com/systems/power/hardware/peripherals/index.html

- Special offers - Power Systems
  http://www.ibm.com/products/specialoffers/us/en/power_systems.html

- *Power ISA Version 2.06 Revision B*
  http://power.org/wp-content/uploads/2012/07/PowerISA_V2.06B_V2_PUBLIC.pdf

- IBM AIX 7.1 Information Center, search for the topics *multiple page size support* and *hardware performance monitor APIs and tools*
  http://pic.dhe.ibm.com/infocenter/aix/v7r1/index.jsp

- The following white papers from Power.org (registration required):

  o *What's New in the Server Environment of Power ISA v2.06*
  https://www.power.org/documentation/whats-new-in-the-server-environment-of-power-isa-v2-06

  o *Commonly Used Metrics For Performance Analysis*
  http://www.power.org/documentation/commonly-used-metrics-for-performance-analysis

  o *Comprehensive PMU Event Reference - POWER7*
  http://www.power.org/documentation/comprehensive-pmu-event-reference-power7

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service. IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law : INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you. This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document was created or updated on March 21, 2013.

Send us your comments in one of the following ways:
- Use the online **Contact us** review form found at:
  ibm.com/redbooks
- Send your comments in an e-mail to:
  redbook@us.ibm.com
- Mail your comments to:
  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400 U.S.A.

This document is available online at http://www.ibm.com/redbooks/abstracts/tips0956.html .

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Active Memory™	POWER Hypervisor™	PowerVM®
AIX®	Power Systems™	PureFlex™
Express Servers™	POWER6®	Redbooks®
IBM®	POWER7®	Redbooks (logo)®
IBM Flex System™	POWER7 Systems™	WebSphere®
POWER®	POWER7+™	
Power Architecture®	PowerLinux™	

The following terms are trademarks of other companies:

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.