

Employing Strategies for Optimizing your IBM POWER8 (and Earlier) Processor-based Systems

IBM Redbooks Solution Guide

Critical business needs can be met with more assurance in an optimized, well-tuned environment. This IBM® Redbooks® Solution Guide describes some of the strategies for optimizing and tuning application code to run on IBM POWER8™ and earlier processor-based systems (IBM Power Systems™). These strategies are drawn from performance optimization efforts across many types of code running on IBM AIX®, IBM i, and Linux®. The guidance ranges from simple to complex and is usable across a broad set of IBM POWER® processor chips and systems. Some of the techniques can be performed without extensive experience. Techniques are also provided for advanced users who have a considerable understanding of application internals. In short, these strategies can maximize the return on your hardware investment.

This technical information was developed by IBM domain experts. It is directed at IBM presales organizations in support of Power Systems products, such as the IBM Power S822 server (Figure 1), and to those responsible for performing migration and implementation activities on IBM POWER8 -based (and earlier) servers.



Figure 1. IBM Power S822 server

Did you know?

Trends in processor design are making it more important than ever to invest in 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 multi-threading (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. IBM Power Architecture supports multiple virtual memory page sizes, which provide performance benefits to an application because of hardware efficiencies that are associated with larger page sizes. POWER8, the newest IBM POWER® processor, supports transactional memory and other advanced features, which can improve application performance.

Business value

You can follow simple strategies and techniques to optimize your POWER8 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 POWER8
- Typically carry over improvements to systems that are based on other POWER processor chips, such as IBM POWER7+™ (and earlier)
- 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 POWER8. Most steps can be carried out without modifying the application source code.
- *Deployment guidelines*, which include steps for configuring POWER8 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 POWER8 processor. The concise introductory guidelines of this Solution Guide and the comprehensive nature of *Performance Optimization and Tuning Techniques for IBM Processors, including IBM POWER8*, SG24-8171, make these valuable resources in your IBM Power Systems environment.

Solution overview

The techniques to optimize your POWER8 environment and to analyze and maximize system performance capitalize on the capabilities and features of the following products:

- The IBM POWER8 processor
- The IBM POWER Hypervisor™
- IBM AIX
- IBM i
- Linux

The IBM POWER8 processor

Several capabilities and features of the POWER8 processor are key to system optimization. POWER8 offers the following 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:

- 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 that are 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.

- Multi-core and multi-thread features, and affinity performance effects

The IBM POWER8 is the latest processor chip in the IBM Power Systems family. The POWER8 processor chip is available in configurations with up to twelve cores per chip, as compared to the IBM POWER7® processor chip, which has up to eight cores per chip. Along with the increased number of cores, the POWER8 processor chip implements SMT8 mode, supporting eight hardware threads per core, as compared to the POWER7, which supported only four hardware threads per core. Each POWER8 processor core supports running in single-threaded mode with one hardware thread, an SMT2 mode with two hardware threads, an SMT4 mode with four hardware threads, or an SMT8 mode with eight hardware threads.

Each SMT hardware thread is represented as a logical processor in AIX, IBM i, or Linux. When the hardware runs in SMT8 mode, the operating system has eight logical processors for each dedicated POWER8 processor core that is assigned to the partition. To gain the full benefit from the throughput improvement of SMT, applications must use all of the SMT threads of the processor cores.

Each POWER8 chip has memory controllers that allow direct access to a portion of the memory DIMMs in the system. Any processor core on any chip in the system can access the memory of the entire system, but 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 POWER8 chip and to the memory that is local to each chip. Software threads that access the same data are best run together on the SMT 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 usage of SMT8 mode, see the "Usage scenarios" section 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, system administrators 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 (AIX 7.1 TL3 SP3 and later) on POWER8 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

With its proven scalability, advanced virtualization, security, manageability, and reliability features, AIX is an enterprise-class OS. In particular, AIX is the only operating system that uses decades of IBM technology innovation to provide the highest level of performance and reliability of any UNIX operating system. Here are some of its benefits:

- Performance benefits
 - Deep integration for Power Architecture
 - Core design with the Power Architecture
 - Autonomic optimization
 - Single OS image configures itself to support any POWER processor
 - Dynamic workload optimization.
 - Performs on a wide variety of system configurations
 - Scales from 0.05 to 256 cores (up to 1024 logical processors)
 - Horizontal (native clustering) and vertical scaling
 - Strong virtualization support for PowerVM virtualization
 - Tight integration with PowerVM
 - Enabler for virtual I/O (VIO)
 - Full set of integrated performance tools

AIX 7.1 runs on and maximizes the capabilities of systems that are based on POWER8, while supporting IBM POWER4® through IBM POWER7+® systems. AIX 6.1 runs on POWER8 systems, but only in POWER7 or IBM POWER6® compatibility modes.

- AIX memory allocation (malloc)

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 multi-thread 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 memory allocation packages that are appropriate for different scenarios. Choose packages by setting environment variables. Select those packages that do not require code modification or application rebuilding. Determining the best malloc package requires an understanding of how applications use the memory allocation routines. To learn how to collect the required information, see Appendix A, "Analyzing malloc usage under AIX", in *Performance Optimization and Tuning Techniques for IBM Processors, including IBM POWER8*, SG24-8171. After the data is collected, experiment with alternatives, alone or in combination.

The following packages are some alternatives that deliver high performance:

- *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 multi-thread applications. The pool malloc is a good choice for both single-thread and multi-thread applications.
- *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 multi-thread applications.

Using the pool front end malloc and the multiheap malloc in combination is a good alternative for multi-thread 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 following environment variable setting:

```
MALLOCOPTIONS=pool,multiheap
```

For more information about using AIX malloc, see the "Usage scenarios" section in this Solution Guide.

- Simultaneous multi-threading (SMT)

Simultaneous multi-threading (SMT) is a feature of Power Architecture and is supported in AIX. AIX provides options to allow SMT customization. The `smtctl` command allows the SMT feature to be enabled, disabled, or capped (SMT2 versus SMT4 mode on POWER7 and SMT2 or SMT4 modes versus SMT8 on POWER8). The partition-wide tuning option, `smtctl`, changes the SMT mode of all processor cores in the partition. It is built on the AIX dynamic reconfiguration (AIX DR) framework to allow hardware threads (logical processors) to be added and removed in a running partition. Because of the global nature of this option, it is normally set by system administrators. Most AIX systems (commercial) use the default SMT settings enabled (that is, SMT2 mode on POWER5 and POWER6, and SMT4 mode on POWER7 and POWER8). When SMT is enabled (SMT2, SMT4, or SMT8 mode), the AIX kernel takes advantage of the platform feature to dynamically change SMT modes. These mode switches are done based on partition load (the number of running or waiting to run software threads) to choose the optimal SMT mode for the CPUs in the partition. The mode switching policies optimize overall workload throughput, but do not attempt to optimize individual software threads.

IBM i

The IBM i operating system and most applications for IBM i are built on a Technology Independent Machine Interface (TIMI) that isolates programs from differences in processor architectures, and allows the system to automatically capitalize on many new Power Architecture features without changes to existing programs. For example, TIMI allows a program to use decimal floating point (DFP) on POWER5 (without special hardware support), and that same program automatically uses hardware support for DFP on POWER6, POWER7, and on POWER8 systems.

IBM Portable Application Solutions Environment for i (PASE for i) is a part of IBM i that allows some AIX application binary files to run on IBM i with little or no changes, so many optimizations that are described for AIX are applicable to PASE for i, including the following ones:

- Multipage size support on IBM i

Most of IBM i uses 4 KB pages, but select system functions automatically use 64 KB pages. Applications running on IBM i 6.1 or later can create shared memory objects that use 64 KB pages (typically using `shmctl` with `SHM_PAGESIZE`). IBM technology for Java programs running on IBM i 6.1 or later can use 64 KB pages for Java heap. PASE for i programs running on IBM i 7.1 or later automatically use 64 KB pages for shared library text and data, and can request 64 KB pages for program text, stack, and data.

IBM Power Systems Firmware does not support 64 KB pages for all configurations. For example, 64 KB pages are not available if a logical partition is configured for IBM Active Memory™ Sharing (AMS).

- Vector Scalar eXtension (VSX)

IBM i 7.2 automatically uses POWER8 vector instructions to improve the performance of some cryptographic operations. PASE for i applications running on IBM i 7.2 on POWER7 or newer processors can use VSX.

Linux

A solid choice for running enterprise-level workloads on POWER8 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. RHEL 6.5 GA and SLES 11 SP3 are the minimum supported versions to run on POWER8 technologies and systems.

IBM Advanced Toolchain V7.0 provides the tuned compilers and libraries for POWER8. The RHEL and SLES 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 POWER8. The typical Linux open source performance tools that Linux users are comfortable with are available on Linux on Power Systems.

Solution architecture

This section describes the architecture of the POWER8 processor and its capabilities for multi-core and multi-thread scalability.

Architecture of the POWER8 processor

The POWER8 processor is manufactured with IBM 22 nm Silicon-On-Insulator (SOI) technology. Each chip is 567 mm² and contains 1.2 billion transistors. The POWER8 processor chip (Figure 2) contains twelve cores. Each core has its own 512 KB L2 and 8 MB L3 (embedded dynamic random access memory (DRAM)) cache, two memory controllers, PCIe Gen3 I/O controllers, and an interconnection system that connects all components within the chip. The interconnect also extends through module and board technology to other POWER8 processors, DDR3 memory, and various I/O devices. The number of memory controllers, memory buffer chips, PCIe lanes, and cores that are available for use depends on the POWER8 system.

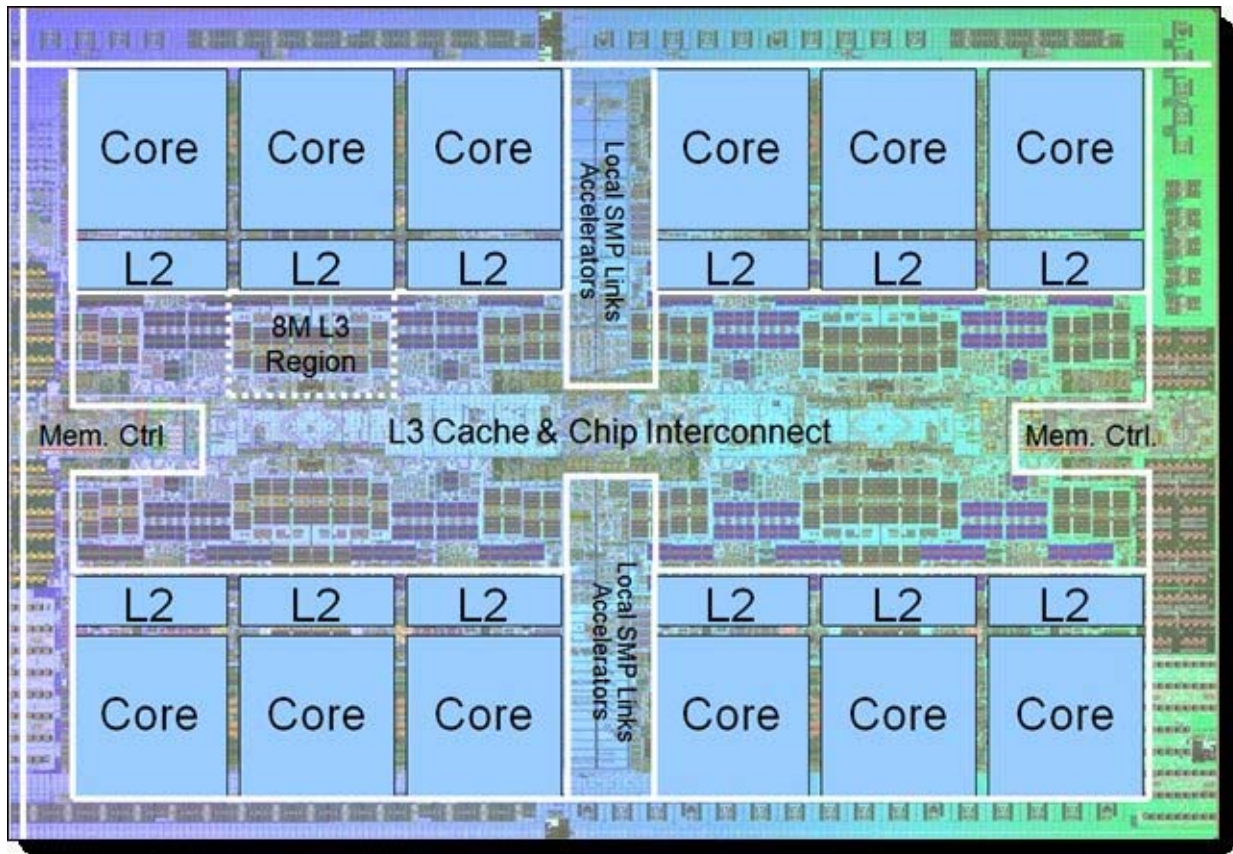


Figure 2. The POWER8 processor chip

Each core is a 64-bit implementation of the IBM Power ISA V2.07 and has the following features:

- Multi-thread design that supports up to an eight-way SMT
- 32 KB, eight-way set-associative L1 i-cache
- 64 KB, eight-way set-associative L1 d-cache
- 72-entry ERAT for effective-to-real address translation for instructions (fully associative)
- 48-entry primary 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
- 16-entry link stack
- 256-entry count cache
- Aggressive out-of-order execution
- Two symmetric fixed-point execution units
- Two symmetric load/store units and two load units, all four of which can also run simple fixed-point instructions

- An integrated, multipipeline vector-scalar floating point unit that supports up to sixteen flops per cycle (eight double precision or sixteen single precision) and that runs the following Scalar and Single Instruction Multiple Data (SIMD)-type instructions:
 - The Vector Multimedia Extension (VMX) instruction set
 - The Vector Scalar Extension (VSX) instruction set
- On-chip encryption
- Hardware data prefetching with 16 independent data streams and software control
- Hardware decimal floating point (DFP) capability

The POWER8 processor is designed for system offerings from single-socket blades to multi-socket enterprise servers. It incorporates a triple-scope broadcast coherence protocol over local and global symmetric multiprocessor (SMP) links to provide superior scaling attributes.

Usage scenarios

This section includes examples of optimization and tuning guidance. For more examples, see *Performance Optimization and Tuning Techniques for IBM Processors, including IBM POWER8*, SG24-8171.

Usage scenario 1: Tuning to capitalize on hardware performance features - SMT8

Some of the best opportunities for improving application performance are in delivering scalable code by having an application make effective use of multiple concurrent threads of execution. With the multi-threaded POWER8 cores, each SMT hardware thread represents a logical processor in AIX, IBM i, and Linux. A 4-core LPAR, with SMT8 mode (eight hardware threads per core), means that the OS is running 32 logical CPUs.

Generally, SMT8 mode is highly valuable for commercial and transactional workloads (such as with IBM WebSphere® Application Server and IBM DB2®), and less valuable for heavy numerical workloads. Occasionally, you must choose between achieving the highest single-thread performance and lowest latency, and achieving the highest system throughput. When it is best to use SMT, here are some of its benefits:

- SMT8 mode is easily enabled on new versions of AIX that support POWER8, where SMT4 is the default. Run `smtctl` to enable, disable, or cap this feature on all processor cores in the partition, and to add or remove hardware threads in a running partition. With SMT8 enabled, your POWER8 system can handle load spikes better and benefit throughput.
- Linux distributions, by default, run at the highest SMT level. Run `ppc64_cpu` to force the system kernel to use lower SMT levels (ST, SMT2, or SMT4 mode), when needed. For example:
 - `ppc64_cpu --smt=1` sets the SMT mode to ST.
 - `ppc64_cpu --smt` shows the current SMT mode.
- IBM i 7.2 adds a job attribute named Processor Resources Priority (PRCRSCPTY) to influence how threads for the job are dispatched. This attribute can request that the system isolate threads for the job on processors that are running fewer threads concurrently, or that the system run threads for the job on processors that are running as many concurrent threads as possible.
- Some earlier versions of operating systems run on POWER8 by using POWER7 compatibility mode. Newer operating systems fully support SMT8 and other new features of POWER8.

Usage scenario 2: 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-threaded application, use the default allocator.
- For a 64-bit application, use the Watson allocator.
- Multi-threaded applications use the `multiheap` option. Set the number of heaps proportional to the number of threads in the application.
- For single-threaded or multi-threaded applications that make frequent allocation and deallocation of memory blocks smaller than 513, use the `malloc pool` 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 size in bucket option) and sizes greater than 512 bytes, use the `configure malloc bucket` option.
- For older applications that require high performance and do not have memory fragmentation issues, use the `malloc 3.1` allocator.
- Ideally, the Watson allocator, along with the `multiheap` and `malloc pool` options, is good for most multi-threaded applications. The pool front end is fast and scalable for small allocations, and, with `multiheap`, ensures scalability for larger and less frequent allocations.
- If you notice high memory usage in the application process even after you run `free()`, the `disclaim` option can help.

For more information, see *Performance Optimization and Tuning Techniques for IBM Processors, including IBM POWER8*, SG24-8171.

Usage scenario 3: Tuning to capitalize on hardware performance features - 64 KB pages

For almost all applications, using 64 KB pages is beneficial for performance. Newer Linux releases (RHEL 5, SLES 11, and RHEL 6) 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 the application must run 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, it is important that the Java heap space uses 64 KB pages, which are enabled by the `-Xlp64k` option in older releases of Java.

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.

For more information, see *Performance Optimization and Tuning Techniques for IBM Processors, including IBM POWER8*, SG24-8171.

Usage scenario 4: 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 POWER8 system that is running in SMT8 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 eight SMT threads, each instance of WebSphere Application Server is bound to 32 logical processors. To ensure good memory and cache affinity on AIX, complete the following steps:

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:
 - o `execrset -c 0-31 -m 0 -e` (command to start the first WebSphere Application Server instance)
 - o `execrset -c 32-63 -m 0 -e` (command to start the second WebSphere Application Server instance)
 - o `execrset -c 64-95 -m 0 -e` (command to start the third WebSphere Application Server instance)
 - o `execrset -c 96-127 -m 0 -e` (command to start the fourth WebSphere Application Server instance)

Consider 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 POWER8 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 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 *Performance Optimization and Tuning Techniques for IBM Processors, including IBM POWER8*, SG24-8171.

Integration

The strategies in this Solution Guide apply to all POWER generations, including the POWER8 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 4 - 48 cores, Power Express servers provide outstanding performance and help 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 POWER8 processor cores with up to 16 TB of memory. It includes the flexibility to turn processors and memory on and off as application workloads dictate.

Linux-only based Power Systems servers

World-class POWER8 systems are equipped with two sockets and up to 32 cores. These value-priced servers go head-to-head with x86 servers in terms of cost and in delivering greater performance, higher usage, 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 of this Solution Guide.

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

Power System model	Part number (feature code)	Charge unit description
IBM Power Systems S812L	8247-21L	The Power System S812L server is designed to provide the ideal foundation for scale-out data and cloud environments in a compact 2U package. This server supports one processor socket, offering ten 3.42 GHz or twelve 3.02 GHz POWER8 cores in a 19-inch rack-mount, 2U (EIA units) drawer configuration. All the cores are active.
IBM Power Systems S814	8286-41A	The Power System S814 server is a powerful 1-socket server that ships with up to eight activated cores. This server supports a one-processor socket, offering 6-core 3.02 GHz or 8-core 3.72 GHz POWER8 processor-based configurations in a 19-inch rack-mount, 4U (EIA units) drawer or desk-side configuration.
IBM Power Systems S822	8284-22A	The Power System S822 server is a powerful 2-socket server that ships with up to 20 activated cores. This server supports two processor sockets, offering 6-core or 12-core 3.89 GHz or 10-core or 20-core 3.42 GHz POWER8 configurations in a 19-inch rack-mount, 2U (EIA units) drawer configuration.
IBM Power Systems S822L	8247-22L	The Power System S822L server is a powerful 2-socket server that ships with 20 or 24 fully activated cores. This server supports two processor sockets offering twenty 3.42 GHz or twenty-four 3.02 GHz POWER8 cores in a 19-inch rack-mount, 2U (EIA units) drawer configuration.
IBM Power Systems S824	8286-42A	The Power System S824 server is a powerful 2-socket server that ships with up to 24 activated cores. This server supports two processor sockets, offering 6-core or 12-core 3.89 GHz, 8-core or 16-core 4.15 GHz, or 24-core 3.52 GHz configurations in a 19-inch rack-mount, 4U (EIA units) drawer configuration.

Related information

For more information, see the following documents:

- *Performance Optimization and Tuning Techniques for IBM Processors, including IBM POWER8*, SG24-8171
<http://www.redbooks.ibm.com/abstracts/sg248171.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:
 - IBM Power Systems S812L (8247-21L)
 - IBM Power Systems S814 (8286-41A)
 - IBM Power System S822 (8284-22A)
 - IBM Power Systems S822L (8247-22L)
 - IBM Power Systems S824 (8286-42A)
- 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/systems.html>
- *Power ISA Version 2.07*
<https://www.power.org/documentation/power-isa-version-2-07>
- 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 paper from Power.org (registration required):
 - *Commonly Used Metrics For Performance Analysis*
<http://www.power.org/documentation/commonly-used-metrics-for-performance-analysis>

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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 grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 actual people or business enterprises is entirely coincidental.

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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

© Copyright International Business Machines Corporation 2014-2017. All rights reserved.

This document was created or updated on April 3, 2017.

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:
redbooks@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/tips1153.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™
AIX®
DB2®
IBM®
IBM Flex System®
POWER®

Power Architecture®
POWER Hypervisor™
Power Systems™
POWER6®
POWER7®
POWER7+™

POWER8™
PowerVM®
PureFlex®
Redbooks®
Redbooks (logo)®
WebSphere®

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.