

New IBM Power systems built with Power10 architecture and processors and Oracle Database



Ravisankar Shanmugam
Wayne T. Martin
IBM Oracle International Competency Center

October 2021



Table of contents

Abstract	1
IBM Power10 processor-based systems	1
IBM AIX and Linux support to IBM Power platform	2
Recommended code levels	3
Service strategy	3
C and C++ compilers	4
Oracle Database and IBM Power servers	4
Capacity comparison of IBM Power10 with Power8 for Oracle Database	6
Tuning tips	6
Oracle Database licensing on IBM Power servers	7
Summary	8
About the authors	9
Document Resources	10
Trademarks and special notices	11



Abstract

This paper provides proof point and interoperability information of the Oracle Database and new IBM Power® systems built with Power10 architecture and processors. Additionally, it provides capacity improvement information for the Oracle Database workloads running on new IBM Power10 processor-based server compared with the previous generation of IBM Power8® processor-based server.

IBM Power10 processor-based systems

On September 8, 2021, IBM announced Power10 processor-based systems, Power E1080 with a combination of one, two, three, or four system nodes per server. Each system node is an enclosure that provides the connections and supporting electronics to connect the processor with the memory, internal disks, adapters, and the interconnects required for expansion. At the time of announcement, a maximum of two system nodes can be ordered, after December 10, 2021, the maximum configuration of four system nodes will be made available. Each system node provides four sockets for Power10 processor chips and 64 differential DIMMS slots for DDR4 technology DIMMS.

Each socket holds one Power10 single chip module (SCM) and contain 10,12, or 15 Power10 processor cores. A 4-node Power E1080 server scales up to 16 processor sockets and 160, 192, or 240 cores and all system nodes must be configured with the same processor feature. Each system node can support up to maximum of 16 TB of system memory and fully configured 4-node Power E1080 can support up to 64 TB of memory.

Each system node provides up to 8 PCIe Gen 5 capable slots, with a maximum of 32 per E1080 server and up to 4 NVMe drive bays per node. Additional drive bays are provided through expansion drawers.

Each system node can have 8 PCIe Gen 5 slots and a maximum of 32 PCIe Gen 5 slots per system. The PCIe Gen 1, Gen 2, Gen 3, Gen 4 adapter cards are also supported in system nodes. Up to 4 PCIe Gen 3 4U I/O expansion drawers per system node providing a maximum of 48 additional PCIe Gen 3 slots and up to 192 PCIe Gen 3 slots using 16 PCIe Gen 3 I/O expansion drawers per system. The system can have over 4,000 directly attached SAS HDDs or SSDs through EXP24SX SFF drawers.

Any nodes system configurations require the system control unit (SCU) to operate. It provides system hardware, firmware and virtualization control through redundant Flexible Service Processors (FSPs).

Each core can run up to eight simultaneous multithreads (SMT) to meet the requirements of resource-hungry applications. The Power E1080 introduces the essential enterprise hybrid cloud platform and uniquely architected to help the customers securely and efficiently scale core operational and AI applications in a hybrid cloud. The Power E1080 simplifies end-to-end encryption and brings AI where the data resides for faster insights.

The Power E1080 server fits a standard 19-inch rack. The server is certified and tested in the IBM Enterprise racks (7965-S42, 7014-T42, 7014-T00, or 7965-94Y), and managed by the Hardware Management Console (HMC) or virtual HMC (vHMC).

For more technical overview, refer IBM Redpaper at the following link:

<https://www.redbooks.ibm.com/redpapers/pdfs/redp5649.pdf>

The following table shows highlights of the new IBM Power10 processor-based system, Power E1080.


New IBM Power10	Number of nodes, cores, speed and I/O slots	Memory (Max)	OS	Server image
E1080	<p>Up to four system nodes.</p> <p>Each node with 4 sockets.</p> <p>Each socket with Power10 Single Chip Module (SCM).</p> <p>Each SCM :</p> <p>10 cores x 3.65 to 3.90 GHz Or</p> <p>12 cores x 3.60 to 4.15 GHz Or</p> <p>15 cores x 3.55 to 4.00 GHz.</p> <p>Integrated PCIe adapter slots :</p> <p>8 PCIe Gen5 per node.</p> <p>PCIe I/O expansion drawers :</p> <p>4 per node with 12 PCIe adapters slots each.</p> <p>Internal Storage :</p> <p>4 NVMe drive bays per node</p>	<p>Each node :</p> <p>64 DIMM slots.</p> <p>Up to 16TB buffered DDR4 CDIMMs.</p> <p>Four nodes:</p> <p>256 DIMM slots.</p> <p>Up to 64 TB buffered DDR4 CDIMMs.</p>	<p>AIX</p> <p>IBM i</p> <p>Linux</p>	 <p>4 System nodes. Each node size is 5U.</p>

Table 1. New IBM Power10 processor-based system, the Power10 E1080

For full specifications of the new IBM Power10 processor-based systems, refer to the following link:

<https://www.ibm.com/it-infrastructure/power/power10>

IBM AIX and Linux support to IBM Power platform

IBM AIX® is an open standards-based UNIX® operating system. AIX, in combination with IBM's virtualization offerings, provide new levels of flexibility and performance to allow you to consolidate workloads on fewer servers, which can increase efficiency and conserve energy. AIX delivers high levels of security, integration, flexibility, scalability and reliability that are essential for meeting the demands of today's information technology environments. AIX operates on IBM systems based on IBM Power Architecture® technology. For more information about AIX, refer to the following web page:

<https://www.ibm.com/it-infrastructure/power/os/aix>

Oracle product support for IBM Power servers is based on AIX operating system support. Oracle does not directly certify IBM server hardware with their software, they certify a specific level of AIX.



To obtain the support information of the AIX version for fully leveraging new features and capabilities of the new IBM Power processor-based servers with physical direct I/O configuration and Virtual I/O, refer the following "System software map for Power servers",

<https://www.ibm.com/support/pages/system-aix-maps>

To obtain the support information for the various Linux® distributions, PowerVM® Virtual I/O Server, IBM i for fully leveraging new features and capabilities of the new IBM Power processor-based servers, refer the following system software map for Power servers:

<https://www.ibm.com/support/pages/node/6020068>

There are two types of Linux distributions available to run on Power E1080:

- Distribution with native Power10 processor technology support
- Distribution running in Power9™ compatibility mode

When an LPAR runs in Power9 compatibility mode, it will benefit from most of the features of the Power10 processor technology, including the full eight threads per core, but program and kernel features using new Power10 instructions or capabilities will not be available.

Distributions with native Power10 processor technology support can also run in Power9 compatibility mode. This is important when doing LPM from a Power9 processor-based server to Power E1080.

The Power9 compatibility mode is required when moving partitions back and forth between Power9 processor based-systems and Power10 processor-based systems. Once a partition has been moved to a Power10 processor-based systems, it can be upgraded to a distribution with native Power10 technology support and rebooted in native Power10 mode.

Please note then is no official support for Power9 compatibility mode for older service packs of SLES 15 such as SLES 15 SP2.

Recommended code levels

There are two tools to help determine the recommended code levels among AIX and Power processors-based server components:

- The Fix Level Recommendation Tool (FLRT) can determine the recommended code levels among a mixture of AIX, HMC, server firmware, VIOS, IBM General Parallel File System (IBM GPFS) and IBM PowerHA®. The FLRT web page is: <http://www14.software.ibm.com/webapp/set2/flrt/home>. Note, the FLRT recommendation provides a minimum acceptable level of compatibility.
- The IBM Power code matrix indicates the recommended code levels for the HMC and server firmware. The Power code matrix web page is: <http://www14.software.ibm.com/webapp/set2/sas/f/power5cm/home.html>. Note that the Power code matrix recommendations can provide the maximum stable code combinations.

Service strategy

To review the latest *IBM AIX Operating System Service Strategy Details and Best Practices* document, refer to the website: <https://www.ibm.com/support/pages/node/883882>



C and C++ compilers

The XL C/C++ compiler family includes several versions of standards-based, high performance C and C++ compilers with advanced optimizing and debugging features. They provide you the ability to optimize and tune applications for optimal execution on systems using all types of IBM Power processors. The compiler family supports IBM Power servers capable of running IBM AIX 7.2, AIX 7.1, and AIX 6.1.

IBM announced IBM Open XL C/C++ for AIX 17.1 on September 8, 2021 that generates code that can take advantage of the capabilities of the latest Power10 architecture and optimize the hardware utilization. It fully supports the IBM Power servers (Power7, Power8, Power9, and Power10) and requires minimum level of AIX 7.2 TL5 SP3.

Customers should carefully evaluate IBM's compatibility requirements when using the available versions of the IBM XL C/C++ compiler to make sure their applications are compatible with the software compiled with the new XL C/C++ compilers and targeted Power platform.

For more information on the binary compatibility of the IBM Open XL C/C++ for AIX 17.1 with previous versions of compilers, refer the following link:

<https://www.ibm.com/docs/en/openxl-c-and-cpp-aix/17.1.0?topic=infrastructure-binary-compatibility>

The IBM XL C/C++ for AIX 16.1 produces code that is compatible with Power10 architecture for both Native mode and Power9 compatible mode. It does not exploit Power10 Native mode.

Open XL C/C++ for AIX 17.1 delivers support and exploitation of the Power10 architecture, support for the C11, C17, C++11, and C++14 programming standards, and implementation of the LLVM compiler infrastructure.

Information about the IBM Open XL C/C++ for AIX 17.1 for IBM Power10 can be seen in the following link:

<https://www.ibm.com/downloads/cas/US-ENUS221-321-CA/name/US-ENUS221-321-CA.PDF>

The C++ applications might have dependency on C++ runtime; and the XL C/C++ runtime environment is installed during the installation of base AIX. To download the latest runtime environment for 16.1 and under, go to the XL C/C++ web page at: <https://www.ibm.com/support/pages/node/795884>

To download the latest runtime environment for 17.1, go to the XL C/C++ web page at:

<https://www.ibm.com/docs/en/openxl-c-and-cpp-aix/17.1.0?topic=reference-open-xl-cc-runtime-environment-filesets>

If the IBM XL C/C++ Enterprise Edition for AIX compiler is installed, confirm that the latest updates are applied by visiting the product support link "[Latest updates for supported IBM C and C++ compilers](#)", then download and apply the appropriate compiler updates.

Oracle Database and IBM Power servers

Oracle Database 19c is the latest long term release and version 19.12 is the latest Release Update (at the time of this publication) for the AIX operating system. Oracle Database 19c includes many new features over its previous database versions. The letter "c" in "19c" stands for "cloud". Oracle Database 19c provides a multitenant architecture that enables Oracle Database to function as a container database



(CDB). A CDB consolidates multiple pluggable databases (PDB), a portable collection of schemas, schema objects, and non-schema objects. The multitenant architecture simplifies the process of consolidating databases into on-premises or cloud models.

Oracle Database 19c allows each database plugged into the multitenant architecture to have the look and feel of a standard database to the applications. Oracle Real Application Clusters (RAC) is an option of Oracle Database that allows a database to be installed across multiple servers (RAC nodes). Oracle RAC uses the shared disk method of clustering databases. Oracle RAC processes running in each node access the same data residing on shared data disk storage.

Oracle Database In-Memory provides a unique dual format architecture that enables tables to be simultaneously represented in memory using traditional row format and an In-Memory column format. The In-Memory Column Store is the key feature of database In-Memory that greatly improves performance of analytics and mixed workloads.

The following table shows the versions of Oracle Database that meet the minimum AIX 7.1 and AIX 7.2 requirement to support on the new IBM Power10 processor-based servers.

Oracle Database	Oracle Database 19c and above Oracle Database 12c Release 2 Oracle Database 12c Release 1 Oracle Database 11g Release 2
-----------------	--

Table 2. Supported Oracle Database versions on the new IBM Power10 server

Even though the above Oracle Database versions are supported with lower versions of AIX 7.1 and 7.2 that are not supported or recommended to use with the new IBM Power10 server, it is recommended to upgrade AIX to the recommended level as mentioned in the above section "IBM AIX and Linux support to IBM Power platform" and the weblink provided for "System software map for Power servers".

For example, the minimum supported AIX 7.2 version for the Oracle Database 19c is AIX 7.2 TL02 SP01. But the minimum supported version AIX 7.2 TL05 SP03 is required for an LPAR created in new Power10 server to use a physical I/O configuration. Oracle Database 19c requires a minimum supported version of AIX 7.2 TL05 SP03 to use a physical I/O configuration in the new IBM Power10 servers.

Oracle Database versions for IBM AIX are upward compatible with the higher versions of AIX Technology Levels (TL) within the same major version. For certification information of Oracle Database with IBM AIX, refer My Oracle Support portal: <https://support.oracle.com/> in the "Certification" section.

Starting with Oracle Database 11g Release 2 (11.2.0.4), Oracle Database Instant Client is supported on Linux on Power (32-Bit) and (64-Bit). Customers currently using the Oracle Database Instant Client 10g Release 2 can upgrade to the 11g Release 2 version by simply replacing the 10g Release 2 libraries with the newer 11g Release 2 version. The Instant Client version of the Oracle Database is also supported on the 19c version of the database. With Oracle Database 12c Release 1, support for little-endian was introduced for the Instant Client running on Linux on Power and with 12c Release 2 support is for little-endian only.



Refer to the *Instant Client Releases Section of My Oracle Support* (see resources) note “*Release Schedule of Current Database Releases (Doc ID 742060.1)*” for the latest status.

Capacity comparison of IBM Power10 with Power8 for Oracle Database

The IBM test team selected two workloads (BI and OLTP type) that are representative of the variety that customers may use. Using Oracle Database 11g Release 2, each workload was driven to high system utilization on a Power8 processor-based system logical partition (LPAR) sized to meet the required throughput objectives. The team then moved the workloads to a Power10 processor-based system LPAR that was sized to provide the capacity similar to Power8 server, based on rPerf ratings from the IBM Power rPerf report found at: <https://www.ibm.com/downloads/cas/K90RQOW8>.

The same database was migrated to a Power10 processor-based system logical partition, and the baseline execution of validation tests were completed on Oracle Database 11g Release 2 and upgraded to Oracle Database 19c (version 19.11). When the workload was again executed, the same high utilization was achieved, and the throughput was recorded.

It has been observed that the movement from Power8 to Power10 provided a benefit of 2.5 times more throughput (delivered by Power10) for OLTP type of workload and improvements delivered in both database versions 11g Release 2 and 19c. Generational improvements in the Power10 architecture increased the number of cores on a chip and improved per core throughput, allowing the partition to be placed on a single system node enclosure.

The BI workload with Oracle Database 19c on IBM Power10 showed 3.61 times improvement on per core basis and 146 times improvement with In-Memory data compared with the result of row type queries on IBM Power8 processor-based systems for the equal number of Power server's rPerf values. The BI workload queries' total elapsed time on Oracle Database version 19.11 was improved up to 2.03 times for queries that ran on row type data and improved up to 79.8 times with In-Memory columnar type data on Power10 processor-based systems compared with Power8 processor-based systems with Oracle Database version 11.2.0.4

For the detailed information about the workloads and results, refer to the following document posted in the at: <https://www.ibm.com/downloads/cas/XOPJMD42>

Tuning tips

The same AIX OS tuning tips are applicable for both Oracle Database 19c and Oracle Database 11g Release 2. For a list of resources that can be useful when tuning an Oracle Database on IBM Power servers, refer to the “Document Resources” section at the end of this paper.

In the IBM AIX 7.2 TL05 SP03, a virtual memory parameter “enhanced_affinity_private” is set to 90 as a default value. The earlier version's default value was 40. If previous versions of AIX (AIX 7.2 TL05 SP02 or lower) are used on a Power10 processor-based system, it is recommended to set the value 90. The purpose of this tuning is to specify the percentage of process private memory allocations that are



affinitized by default. This tunable parameter limits the default amount of affinitized memory allocated for process private memory. Affinitizing private memory may improve process throughput. The result may vary for different types of workloads.

If using Oracle Database 19c In-Memory feature, the queries can take advantage of the tuning parameter `INMEMORY_OPTIMIZED_ARITHMETIC=ENABLE`, for tables compressed with "QUERY LOW" option, the NUMBER data type encoded as a fixed-width native integer scaled by a common exponent. In general, the NUMBER can incur a significant performance overhead for queries because the arithmetic operations cannot be performed natively in hardware. The In-Memory optimized number format enables native calculations to happen in hardware for the database segments compressed with the "QUERY LOW" (default) compression option.

Oracle Database licensing on IBM Power servers

Oracle's licensing policies can be found at Oracle's "Oracle Global Pricing and Licensing" website at: <https://www.oracle.com/corporate/pricing/index.html>

In general, Oracle's licensing policy applies equally to any hardware platform with a few exceptions. For the purpose of this white paper, the Core Factor table is a critical document since Oracle applies different core factors to different hardware platforms in order to calculate the price of its software per core.

At the time of this writing Oracle has not published the core factor for the IBM Power10 processor-based server, therefore the value is one based on the "all other Multi core chips" category.

In some cases, Oracle software can be licensed by number of users instead of licensing per core. The Oracle core factor table is available at:

<http://www.oracle.com/us/corporate/contracts/processor-core-factor-table-070634.pdf>.

The new IBM Power10 processor-based servers are all configured with up to 4 system nodes, each node has 4 sockets, where each socket is populated with an SCM. One system node can be configured with up to 60 cores in the 4 sockets.

Three versions of the Oracle Database are available for use on IBM Power servers:

- **Oracle Database Standard Edition** is only applicable on the servers with Oracle Database 11g Release 2 based on current Oracle product support. It can be used on a system that has a maximum of 4 sockets.
- **Oracle Database Standard Edition 2** is supported on systems or clusters of servers with a maximum of two sockets. The Power10 E1080 server is not eligible to run with Oracle Database Standard Edition 2.
- **Oracle Database Enterprise Edition** is licensed by the number of physical cores available in the logical partitions that the Oracle Database is running in. This is because Oracle accepts Dynamic Logical Partitions (DLPAR) as hard partitioning and charges for the maximum number of physical cores that the Oracle Database is capped to run on. Any IBM Power server can run Oracle Database Enterprise Edition.

IBM and IBM Business Partners can contact the IBM Oracle International Competency Center (ibmoracle@us.ibm.com) for help with evaluating the number of cores available in an LPAR.



The document at the following link shows Oracle's partitioning policies:

<http://www.oracle.com/us/corporate/pricing/partitioning-070609.pdf>

Oracle core factors are applied to licenses of Oracle Database Enterprise Edition. To calculate the number of Oracle Database Enterprise Edition licenses required the core factor is applied to the total number of cores the database will run on and the factored value is used to calculate the number of licenses required. Oracle uses the core factor to adjust for the performance of the processor with more powerful processors having a higher core factor. However, this determination is exclusively the responsibility of Oracle Corporation and a client planning to install Oracle software on any hardware platform needs to discuss licensing terms and conditions with their Oracle representative.

Summary

This paper gathers the key technical topics that relate to the newly introduced IBM Power10 processor-based server the Power E1080 to help understanding of the support of AIX versions, Oracle Database versions, and IBM virtualization software versions. Additionally, a set of capacity improvement results were included in this document from capacity improvement evaluation studies with a Power E1080 server with OLTP and analytical workloads running on Oracle Database 19c and Oracle Database 11g Release 2.

The capacity improvement numbers shown in this document are for education purpose only. The results were derived from configurations that used default values and generally accepted best practices without any intense tuning of AIX, Oracle Database 11g Release 2, Oracle Database 19c, or the storage area network (SAN) server. The results would vary on different types of applications with differing workload characteristics.



About the authors

Ravisankar Shanmugam is an Advisory Software Engineer with the ISV Technical Enablement team in the IBM Systems group based in Foster City, California. He writes technical collateral, leads proof of concept and performance evaluation projects for IBM Power servers with Oracle products. He can be reached at ravisankar@us.ibm.com.

Wayne Martin is the IBM Systems Technology Solutions Manager responsible for the technology relationship between IBM and the developers of Oracle Corporation's Database and Fusion Middleware for all IBM server brands. His responsibilities include driving the mutual understanding between IBM and Oracle on technology innovations that will generate benefits for mutual customers, managing the process of getting that technology implemented in products and ensuring that availability of the products to customers is timely. Wayne has held a variety of technical and management roles at IBM that have focused on driving enhancements of ISV software that uses IBM's mainframe, workstation, and scalable parallel products. He can be reached at wmartin@us.ibm.com.



Document Resources

The following are some useful document resources:

- Modernizing Oracle Database on IBM Power:
<https://www.ibm.com/downloads/cas/X0PJMD42>
- System Software Maps:
<https://www.ibm.com/support/pages/system-software-maps>
- Service and support best practices for Power Systems:
<https://www.ibm.com/support/pages/node/883882>
- IBM FlashSystem 9100 Architecture, Performance, and Implementation:
<https://www.redbooks.ibm.com/redbooks/pdfs/sg248425.pdf>
- Managing the Stability and Performance of current Oracle Database versions running AIX on Power Systems including POWER9:
<https://www.ibm.com/support/pages/node/6355543>
- Oracle Database 19c & Oracle Database 19c RAC on IBM AIX Tips and Considerations:
https://www.ibm.com/support/pages/system/files/inline-files/IBM_AIX_Oracle_19c-tips_Shanmugam_Oct2019.pdf
- IBM Open XL C/C++ for AIX:
<https://www.ibm.com/products/open-xl-cpp-aix-compiler-power/details>

These are some supplemental tuning resources available at the My Oracle Support web site (userid and password are required to sign in):

- <https://support.oracle.com/>



Trademarks and special notices

© Copyright. IBM Corporation 2021. All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

AIX, IBM, the IBM logo, Power, Power Architecture, PowerHA, PowerVM, POWER8, and POWER9 are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

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

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

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

The information provided in this document is distributed "AS IS" without any warranty, either express or implied.

The information in this document may include technical inaccuracies or typographical errors.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.



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.



Please recycle