Modernize Oracle Database infrastructure from IBM Power9 to Power11

Evaluate capacity, efficiency, and licensing implications for Oracle Database workloads

Table of Contents

Overview 3

Disclaimer 3

Software and hardware configuration 4

Oracle Database and IBM Power configurations 4

Brokerage – OLTP workload 5

- Partition placement in Power servers for OLTP workload 7
- Capacity comparison of Power11 with Power9 running OLTP workload 7

BIDAY analytical workload 10

- Partition placement in Power servers for BIDAY queries 12
- Capacity comparison of Power11 with Power9 running BIDAY queries 13

CPU utilization and I/O throughput 15

Oracle Database In-Memory (IM) 15

Conclusion 17

Appendix 17

Get more information 19

About the authors 19

Acknowledgement 20

Overview

The challenge

Customers need to upgrade IBM® Power® for Oracle® Database to ensure greater capacity and optimize total cost of ownership (TCO), including processor core licensing and energy consumption.

The solution

IBM Power11™ is the ideal choice for organizations running Oracle Database, offering improved capacity, scalability, security, and reliability, while enabling the adoption of current and future ready technologies in a cost-effective manner.

The objective of this white paper is to demonstrate the advantages of IBM's latest-generation IBM Power11 processor-based servers over previous generations of IBM Power servers, such as IBM Power9® processor-based systems, for Oracle Database workloads. The Power11 platform delivers substantial improvements in per-core performance, scalability, and energy efficiency, enabling enterprises to achieve higher throughput and faster response times with fewer Power11 cores compared with previous-generation Power servers for Oracle database workloads. By leveraging these advancements, organizations can optimize resource utilization, reduce licensing and energy costs, lower TCO, and support demanding mission-critical Oracle Database environments with improved efficiency, security, and reliability.

In this sample capacity evaluation test, a logical partition (LPAR) created on an IBM Power9 processor-based server (Power E980) and another logical partition created on an IBM Power11 processor-based server (Power E1180) were used. The IBM test team used two internally developed workloads: one based on online transaction processing (OLTP) and the other on analytical processing, both representative of the variety that enterprises may use. Each workload was run on the Power9 processor-based server with Oracle Database 19c, and the same version of the database was deployed on the Power11 processor-based server.

Key findings from the technical evaluation: IBM Power11 processor-based servers demonstrate significant per-core performance and efficiency improvements over Power9 servers for Oracle Database workloads, enabling higher throughput and lower response times using fewer processor cores. These gains can help reduce Oracle Database licensing requirements, lower power consumption, and improve overall infrastructure efficiency for mission-critical database environments.

Disclaimer

The results shown in this paper are for education purposes only. The results do not represent the full potential capability of Power11 and Power9 processor-based systems. The results were derived from configurations that used default values and generally accepted best practices, without any intense tuning on IBM AIX®, Oracle Database 19c, or the storage area network (SAN) storage disk drives. The results may vary on different Power11 and Power9 processor-based systems and for different types of applications with differing workload characteristics.

Software and hardware configuration

This section describes the software and hardware configurations used for the capacity evaluation. These details establish the technical context required to interpret the workload behavior and performance results presented later in this paper.

Software

The software stack listed below defines the operating system, database, and supporting components used consistently across both Power9 and Power11 test environments.

- Oracle Database 19c (19.27/19.28)
- IBM AIX 7.2 TL05 SP10
- IBM AIX 7.3 TL03 SP01

Hardware

The hardware specifications listed below describe the IBM Power server models and system characteristics used for executing the Oracle Database workloads in this evaluation.

- Power11: IBM Power E1180
- Power9: IBM Power E980

Oracle Database and IBM Power configurations

This section describes the Oracle Database and IBM Power server configurations used for the capacity evaluation. The configuration details provide the baseline technical context for understanding the workload execution and capacity comparisons presented in the following sections.

For this sample experiment, a logical partition was created on an IBM Power E980 server running AIX 7.2 TL05 SP10, with dedicated processor cores, Fibre Channel (FC) adapters, and SAN volumes assigned from the IBM FlashSystem® 7300. Oracle Database 19c (single instance) was then installed and configured for the workloads. The Oracle Database workloads were executed under different load levels and configurations.

After completing the workload runs, all volumes used for AIX and the Oracle Database were moved to the logical partition created on the IBM Power E1180 server. The logical partition on the Power11 server was configured with dedicated processor cores and Fibre Channel adapters, mapped to the same SAN volumes previously used by the Power9 server logical partition. AIX was migrated from version 7.2 TL05 SP10 to version 7.3 TL03 SP01, and the same Oracle Database workloads were executed under the same test scenarios used on the Power9 server.

The workloads were driven to high system utilization on both Power9 and Power11 logical partitions to meet required throughput and response time objectives. The logical partitions were sized with physical cores to provide the same capacity based on rPerf ratings from the IBM Power Performance Report.

The rPerf is an IBM benchmark metric used to estimate and compare the performance of IBM Power systems running AIX. It measures a system's relative throughput capacity, helping with capacity planning, sizing, and performance comparison across Power processor generations. The scale is based on a single Power4 processor (rPerf = 1.0), with higher values indicating greater performance (for example, 10.0 = ten times the baseline). Derived from IBM's internal simulations of enterprise workloads such as OLTP and batch processing, rPerf reflects architectural improvements in newer processors (for example, Power11 versus Power9). However, it is an estimate under ideal conditions and may differ from the real-world performance.

Subsequent sections of this paper provide detailed configuration information for the Power9 and Power11 logical partitions, and the Oracle Database workloads executed during the tests.

The evaluation begins with an OLTP workload, as OLTP environments are commonly capacity-constrained and are sensitive to per-core performance, making them an effective starting point for comparing processor generations.

Brokerage – OLTP workload

This section describes the OLTP workload used to evaluate processor-generation differences between IBM Power9 and Power11 servers under transaction-intensive conditions.

The brokerage OLTP workload simulated the transactions and database activity of a stock brokerage firm. This workload performs primarily read operations with fewer write operations. The test team created a database of approximately 1 TB in size containing 50,000 customer records. The team modified the transaction mix and tuned the database system global area (SGA) size so that the read/write ratio of the workload is approximately 90/10 and more central procession unit (CPU) operations are performed. Table 1 provides the hardware and software configuration details (as returned by the `prtconf` command). This configuration was used for the sample brokerage OLTP workload. The Oracle Database initialization parameters used for executing the OLTP workload are provided in the Appendix.

The Ansible® collection **power_aix_oracle** was used to install Oracle Database 19c, version 19.28. This collection simplifies and automates the installation of Oracle Database 19c on the AIX operating system. For more details, refer to the following links.

- Ansible Galaxy - IBM Power AIX Oracle
- GitHub - IBM Ansible Power AIX Oracle

Configuration	Power9	Power11
System model	IBM, 9080-M9S	IBM, 9080-HEU
Processor type	PowerPC_Power9	PowerPC_Power11
Number of cores	24	16
Clock speed	2.9 GHz to 3.8 GHz	3.8G Hz to 4.2 GHz
Memory	800 GB	800 GB
Firmware version	IBM, FW950.D0 (VH950_175)	IBM, FW1110.00 (RK1110_060)
Nodes/drawers	1	1
Chips per node	4	1
simultaneous multithreading (SMT) (default values used)	8	8
OS level	AIX 72 TL5 SP10	AIX 73 TL3 SP1
Oracle Grid Home Version	19c (version 19.28)	19c (version 19.28)
Oracle DB home version	19c (version 19.28)	19c (version 19.28)
Workload users	216	216

Table 1: Hardware and software configuration

Figure 1 shows the migration methodology that was followed for the brokerage OLTP workload across the Power9 and Power11 logical partitions during the evaluation.

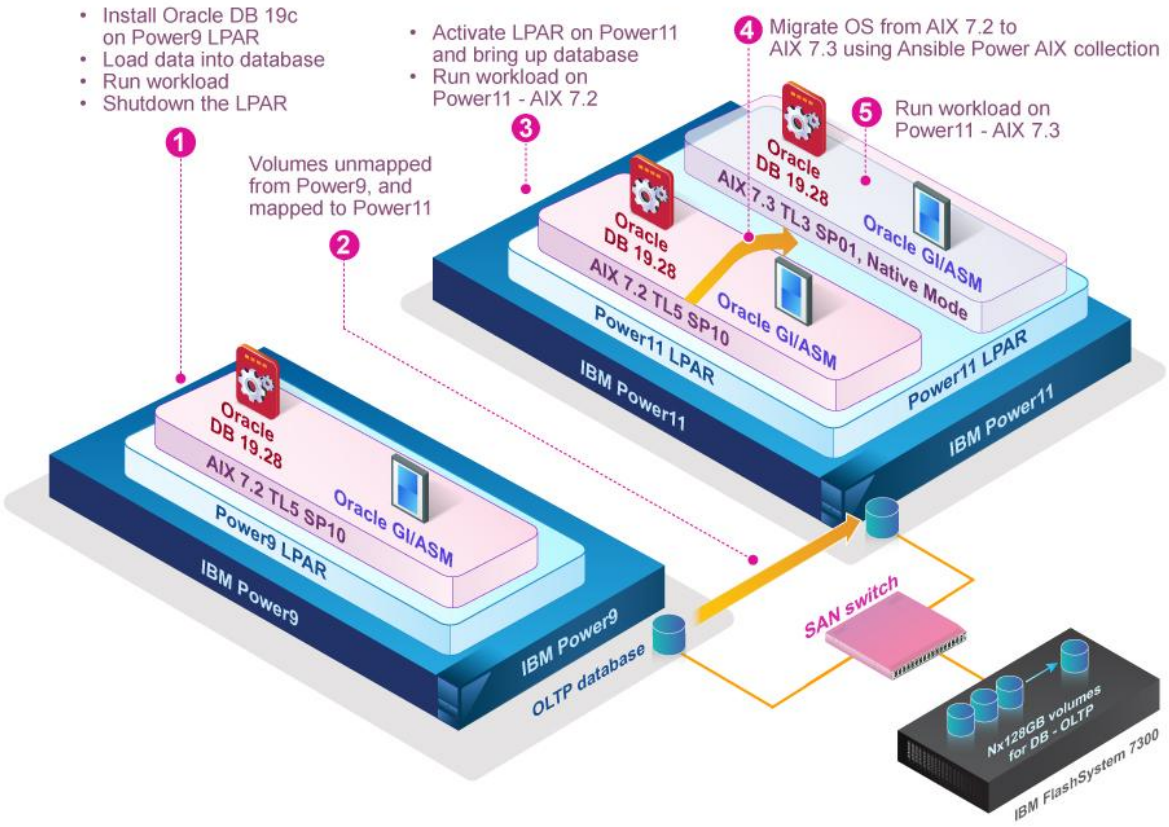


Figure 1: Flow diagram for OLTP workload

As shown in the flow diagram in Figure 1, capacity metrics were collected during the Oracle Database OLTP workload execution on the Power9 LPAR. After data collection, the database was stopped, and the AIX and database volumes were moved to the Power11 LPAR. The database was then started on the Power11 LPAR to collect baseline metrics for the OLTP workload. After completing the baseline validation tests on AIX 7.2, the operating system was migrated to AIX 7.3, and workload execution was resumed.

Partition placement in Power servers for OLTP workload

The `lssrad` command was used to obtain the partition-placement details. On the Power11 server, a 16-core LPAR was created and resized to fit within a single socket (chip), as shown in Figure 2.

```
REF1   SRAD        MEM      CPU
0
          0   795491.25      0-127
1
          1       0.00
```

Figure 2: Output of the `lssrad` *command on the Power11 server LPAR*

On the Power9 server, a 24-core LPAR was configured, with the processors and memory fitting within a single node enclosure, as shown in Figure 3.

```
REF1   SRAD        MEM      CPU
0
          0   198907.38      0-47
          1   198702.00      48-95
          2   198935.00      96-143
          3   198933.44      144-191
1
          4       0.00
```

Figure 3: Output of the `lssrad` *command on the Power9 server LPAR*

Capacity comparison of Power11 with Power9 running OLTP workload

Based on rPerf values that were published for Power9 and Power11, it was observed that the LPAR in a 24-core Power9 processor-based system is equivalent to a 16-core Power11 processor-based system. The test team ran the same number of workload users on Power9 24-core LPAR and on Power11 16-core LPAR, and CPU utilization reached approximately 100%. Initially, the test performed the OLTP workload baseline runs on the Power9 24-core LPAR having the operating system AIX 7.2 TL5 SP10. Next, the Power9 LPAR was reduced, and all disks were unmapped from Power9 and mapped to Power11. The 16-core Power11 LPAR was activated, and the OLTP workload run was performed on the AIX 7.2 TL5 SP10 operating system. Next, on Power11 LPAR, the AIX operating system was migrated from AIX 7.2 TL5 SP10 to AIX 7.3 TL3 SP1 using the Ansible collection power_aix playbook demo_nim_alt_disk_migration.yml.

Using the same workload mix, the team performed the runs on the Power11 LPAR having the operating system AIX 7.3 TL3 SP1. Figure 4 illustrates the relative throughput per core, and Figure 5 illustrates the relative response time per core improvement.

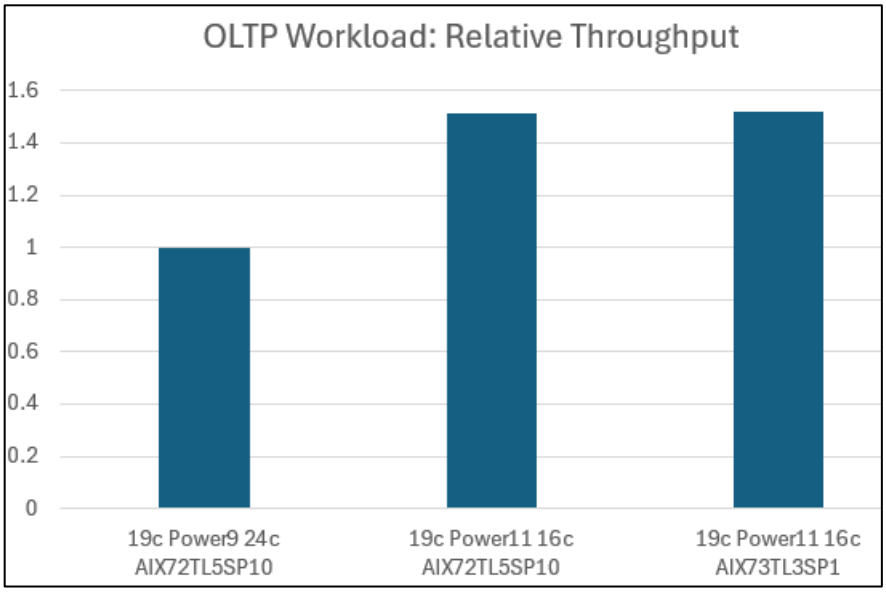


Figure 4: OLTP workload relative throughput time per core

The following formula was used to calculate the per-core response time improvement factor for Power9 and Power11 servers.

$$\text{Relative response time per-core improvement factor ratio} = \frac{\text{Power9 cores used} \times \text{average response time in Power9 LPAR}}{\text{Power11 cores used} \times \text{average response time in Power11 LPAR}}$$

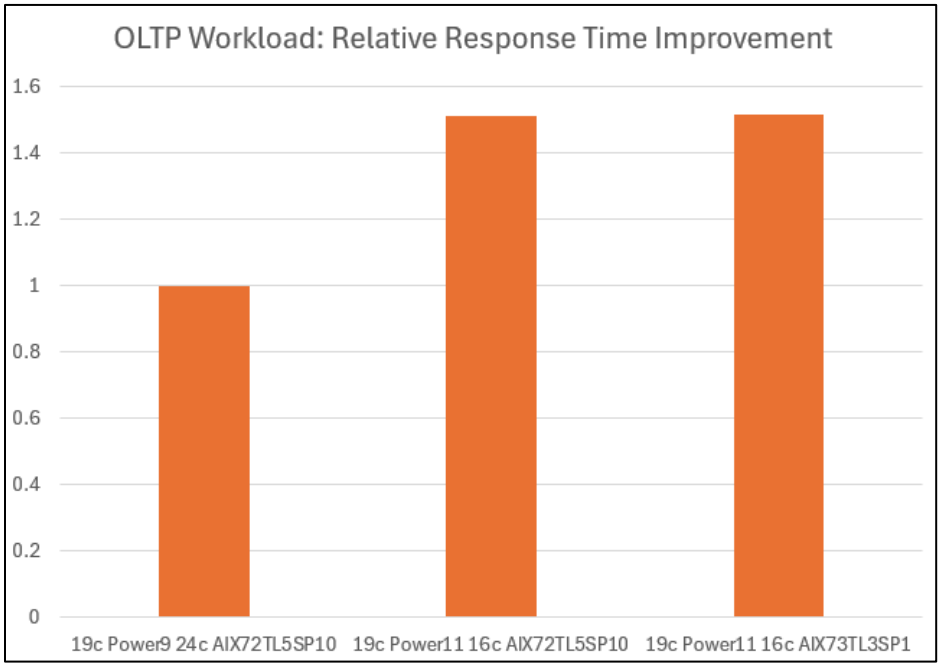


Figure 5: OLTP workload relative response time per-core improvement factor

The following observations can be made from Figures 4 and 5:

- The execution plans (`plan_hash_value`) for the workload Structured Query Language (SQL) statements did not change when the test team moved the storage volumes from the Power9 to the Power11 system. Throughput improvements were observed after restarting the Oracle Database 19c on the Power11 processor-based system.
- After migrating the operating system from AIX 7.2 TL05 SP10 to AIX 7.3 TL03 SP1, the execution plans for the workload SQL statements remained unchanged, with a small improvement in the observed throughput.
- When migrating from Power9 to Power11, the percentage improvement in throughput was found to be directly proportional to the percentage improvement in response time.

It was observed that the movement from Power9 to Power11 provided a benefit of approximately 1.5 times more throughput (delivered by Power11). Generational improvements in the Power11 architecture increased the number of cores on a chip and improved the per-core throughput, allowing the partition to be placed on a single socket (chip). This improved throughput and reduced the space and power requirements for the workload.

Having evaluated transactional performance characteristics using the OLTP workload, the evaluation next focuses on analytical processing behavior using the BIDAY workload, which represents CPU- and I/O-intensive business intelligence query patterns.

BIDAY analytical workload

For the sample test experiment running analytical workloads on the IBM Power9 and Power11 LPARs, an IBM-developed synthetic analytical workload, BIDAY, was used. BIDAY models various types of analyses commonly performed in business intelligence (BI) workloads. The workload has about 1 TB of raw data representing retail sales tracking loaded into Oracle Database.

The BIDAY workload consists of a set of 26 queries with complexity ranging from simple to very complex queries. The workload can be scaled by executing 1, 2, 4, 8, and more concurrent users. The BIDAY schema includes a table named **sales_fact**, which is a fact table that is single range-hash partitioned and contains nine billion rows.

During the test, the set of all 26 queries was run sequentially in the same strict order, and the team presented each concurrent execution of this set as an additional user. For example, for two users two copies of each query set were run, and for eight users eight copies of each query set were run concurrently. The execution of the queries was serialized, and therefore, all the users executed the same query set.

The configuration (returned by using the `prtconf` command) shown in Table 2 was used with an IBM AIX LPAR for evaluating the capacity of the Oracle Database workload on IBM Power servers. The Oracle Database initialization parameters used for executing the BIDAY workload are provided in the Appendix.

LPAR configuration	IBM Power9 LPAR	IBM Power11 LPAR
System model	IBM 9080-M9S	IBM 9080-HEU
Processor type	PowerPC_Power9	PowerPC_Power11
Number of cores used	12 to 32 (dedicated)	6 to 24 (dedicated)
Clock speed or workload optimized frequency (WOF) range	2.9 GHz to 3.8 GHz	3.8 GHz to 4.2 GHz
Memory	512 GB	512 GB
Firmware version	FW950.D0 (VH950_175)	FW1110.00 (RK1110_060)
Nodes/drawers	1	1
Chips per node	1 to 4	1 to 4
SMT (default values used)	8	8
OS level	AIX 7200-05-10	AIX 7300-03-01
Oracle Grid Infrastructure	19c version 19.27	19c version 19.27
Oracle Database	19c version 19.27	19c version 19.27
Internal storage adapters	2 x PCIe3 32 Gb FC ports	2 x PCI4 32 Gb FC ports
SAN disk drives	16 x 128 GB SAN external storage logical unit numbers (LUNs)	16 x 128 GB SAN external storage LUNs

Table 2: Configuration used with an IBM AIX LPAR

In Figure 6, the LPARs are defined with their own profiles on their respective servers as shown in Table 3. The profiles created on the Power9 and Power11 servers each include two 32 GB Fibre Channel ports, several dedicated processors based on their server type and test scenarios (either Power9 or Power11 cores), as well as assigned memory and network adapters. The SAN storage LUNs are mapped through the Fibre Channel ports allocated to these profiles.

Initially, the LPAR on the Power9 server was activated with IBM AIX 7.2 TL05 SP10 installed on a SAN LUN. Then Oracle Database 19c (version 19.27) was then installed, and an Automatic Storage Management (ASM) disk group with external redundancy was created. The BIDAY raw data was subsequently imported into the pluggable database container.

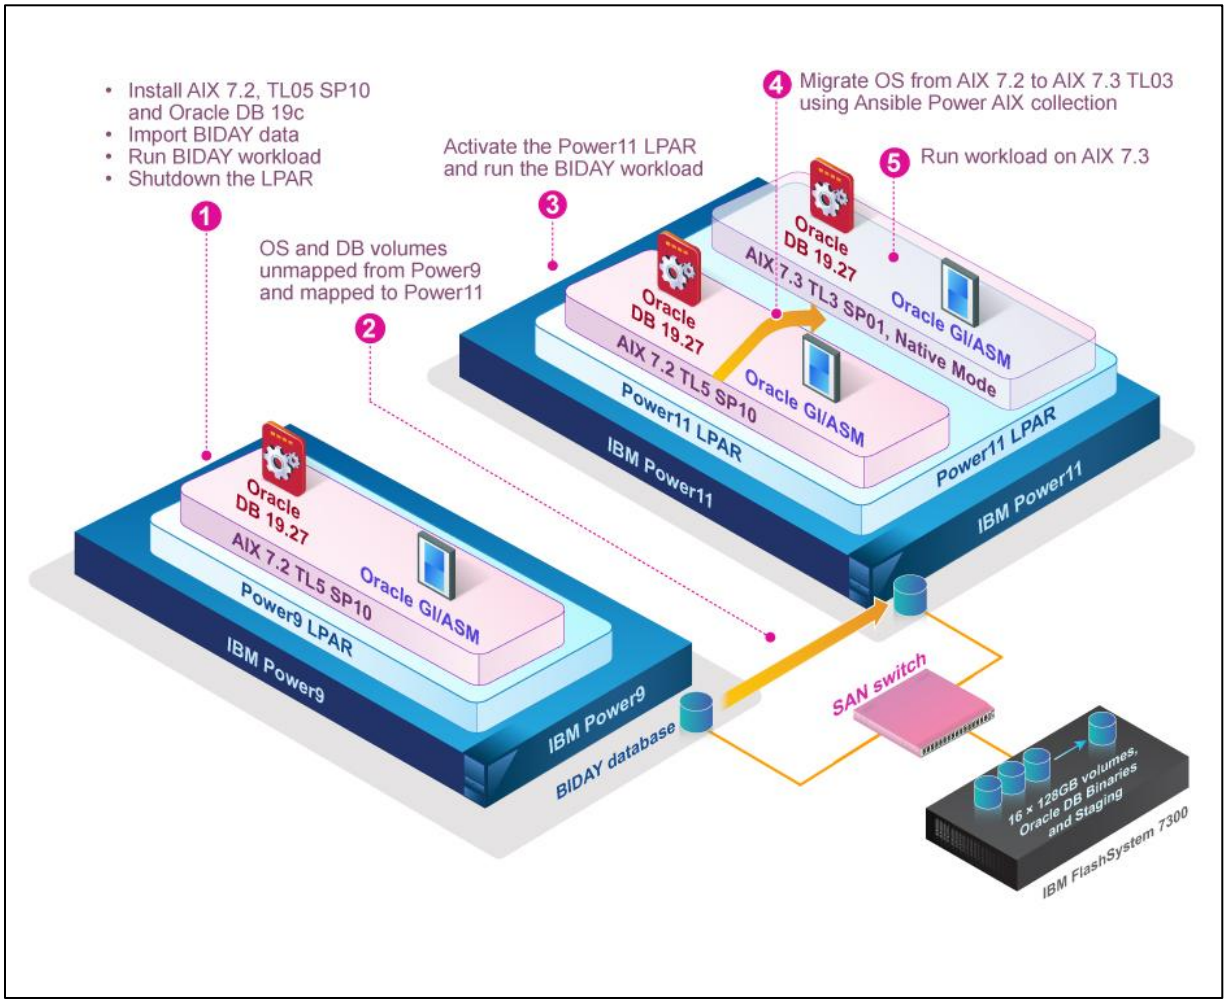


Figure 6: Flow diagram for BIDAY workload, Oracle Database, and AIX 7.x

As shown in the flow diagram in Figure 6, capacity metrics were collected during the BIDAY workload execution on the Power9 LPAR. After data collection, the database was stopped, and the AIX and database volumes were moved to the Power11 LPAR. The database was then started on the Power11 LPAR to collect baseline metrics for the BIDAY workload. After completing the baseline validation tests on AIX 7.2, the operating system was migrated to AIX 7.3 TL03 SP01. The LPAR was then restarted to enable Power11 native mode, and BIDAY workload execution was resumed.

The following section describes how processor cores and memory were placed on the Power9 and Power11 systems to ensure comparable resource alignment for the BIDAY workload and to provide context for the capacity improvement results that follow.

Partition placement in Power servers for BIDAY queries
In the IBM Power E1180 (Power11, model 9080-HEU), the server uses a single chip module (SCM) design, meaning each socket contains a single processor chip. Each of these SCMs can be configured with 10, 12, or 16 cores per socket, depending on performance and power requirements. The system is built to scale, supporting up to four system nodes, each with up to four sockets, for a maximum of 16 sockets per system. In its highest configuration, the IBM Power E1180 provides up to 256 processor cores (16 cores × 16 sockets). The IBM Power E1180 system used for this experiment was fully configured with all 256 cores.

The processor cores and memory resources assigned to the LPAR were aligned to be placed in a single system node enclosure in both the Power11 processor-based server and the Power9 processor-based server. The assigned memory was almost equally shared across the cores.

On the Power11 LPAR, the default assignment of 24 dedicated cores with SMT8 and 512 GB of memory can be viewed using the following IBM AIX® command.

Figures 7 through 10 illustrate the allocation and placement of processor cores and memory for different Power9 and Power11 LPAR configurations used during the BIDAY workload evaluation.

```
#lssrad -av
REF1   SRAD        MEM      CPU
0
          0   127450.75      0-47
          1   127239.00      48-95
          2   127239.00      96-143
          3   127216.00      144-191
```

Figure 7: LPAR cores and memory placement in Power11 server for 24 dedicated cores and 512 GB memory

The assignment of 16 cores and 512 GB of memory to the Power11 server LPAR is shown in the following output.

```
# lssrad -av
REF1   SRAD        MEM      CPU
0
          0   509125.25      0-127
```

Figure 8: LPAR cores and memory placement in Power11 server for 16 dedicated cores and 512 GB memory

In the Power9 server LPAR, the assignment of 32 dedicated cores with SMT8 and 512 GB of memory is shown in the following output.

```
# lssrad -av
REF1   SRAD        MEM      CPU
0
          0   127432.19      0-63
          1   127212.00      64-127
          2   127239.00      128-191
          3   127239.00      192-255
```

Figure 9: LPAR cores and memory placement in Power9 server for 32 dedicated cores and 512 GB memory

The assignment of 12 cores with SMT8 and 512 GB of memory to the Power9 server LPAR is shown in the following output.

```
#lssrad -av
REF1   SRAD        MEM      CPU
0
          0   509125.75      0-95
```

Figure 10: LPAR cores and memory placement in Power9 server for 12 dedicated cores and 512 GB memory

The BIDAY workload queries ran on Oracle Database 19c version 19.27.0.0 with 1, 2, 4, and 8 users with the `parallel_degree_limit DB` parameter set to a number that matched the number of logical processors available in the LPAR. The `parallel_degree_policy` parameter was set to AUTO. For example, the LPAR in the IBM Power11 processor-based system was set up with 24 dedicated cores with SMT8 enabled. In total, 192 logical processors are available to the Oracle Database for executing the queries with up to 192 parallel processes.

Capacity comparison of Power11 with Power9 running BIDAY queries
The IBM Power11 LPAR with 24 dedicated cores outperformed the Power9 LPAR with 32 dedicated cores for running the BIDAY workload on the data residing in the database in the traditional row method.

To ensure a fair comparison between Power9 and Power11 systems, the effective capacity of processor cores was normalized using IBM Power rPerf values, as described in the following calculation and comparison.

The number of Power11 and Power9 cores is determined based on the rPerf values of the respective Power server types. Servers with older-generation Power processors, such as Power9, have lower rPerf (relative performance) values compared to servers with newer-generation processors, such as Power11. Therefore, 32 dedicated Power9 cores are equivalent in capacity to 24 dedicated Power11 cores, based on the measurements of the Power9 and Power11 servers used in this test.

The number of Power11 cores equivalent to Power9 cores is calculated as follows:

Power9 server
9080-M9S-SN7835D68 (E980) with a total of 144 cores and an rPerf value of 3810.6 (SMT8)
1 Power9 core = 3810.6 ÷ 144 = **26.46 units**

Power11 server
9080-HEU-SN78BBB08 (E1180) with a total of 256 cores and an rPerf value of 9141 (SMT8)
1 Power11 core = 9141 ÷ 256 = **35.7 units** (corrected from 37.71; 9141 ÷ 256 = 35.7)

From these values:
32 Power9 cores = 32 × 26.46 = 846.72 units

Equivalent number of Power11 cores = 846.72 ÷ 35.7 = 23.71 cores (rounded to **24 Power11 cores**)

To compare effective per-core efficiency across Power9 and Power11 systems while accounting for both execution time and concurrency, the following derived metrics were used.

While comparing the effective capacity across scenarios, a metric that accounts for both the improvement in core capacity and the reduction in query elapsed time is required. For this purpose, the test team used a metric calculated as follows:

$$X = \frac{(\text{Power9 cores used} \times \text{query elapsed time (seconds) in Power9 LPAR})}{\text{Number of users}}$$

$$Y = \frac{(\text{Power11 cores used} \times \text{query elapsed time (seconds) in Power11 LPAR})}{\text{Number of users}}$$

$$\textbf{Relative per-core improvement factor ratio} = \frac{\mathbf{X}}{\mathbf{Y}}$$

The Power11 system was observed to improve the BIDAY queries by up to average 2.26 to 2.33 times compared with the Power9 system.

Capacity metrics: Per-core improvement in the execution of queries in row format	BIDAY workload with number of concurrent users			
	1 user	**2 users**	**4 users**	**8 users**
Relative per core capacity improvement factor ratio with **24** Power11 cores over **32** Power9 cores	2.28x	2.32x	2.22x	2.22x
Relative per core capacity improvement factor ratio with **16** Power11 cores over **32** Power9 cores	2.42x	2.34x	2.27x	2.30x

Table 3: Per core capacity comparison of Power11 with Power9 for BIDAY row method

Even when the number of Power11 cores is reduced to half the number of Power9 cores, the Power11 cores still outperform the Power9 cores.

While the per-core improvement ratios illustrate relative efficiency across processor generations, the following table presents the corresponding percentage improvement in query response time for the same workload scenarios.

The following table shows the percentage improvement in capacity metrics on Power11 compared to Power9. The formula used for calculating the percentage improvement is as follows:

For each user, **capacity increase (%)** =

$$\frac{(\text{Elapsed time of queries in Power9 LPAR} - \text{Elapsed time of queries in Power11 LPAR})}{\text{Elapsed time of queries in Power9 LPAR}} \times 100$$

Capacity metrics: Percentage increase in the execution of queries in row format	BIDAY workload with number of concurrent users			
	1 user	**2 users**	**4 users**	**8 users**
Relative capacity increase percentage with **24** Power11 cores over **32** Power9 cores	41.47%	42.58%	40.05%	39.87%
Relative capacity increase percentage with **16** Power11 cores over **32** Power9 cores	17.51%	14.66%	11.78%	13.11%

Table 4: Capacity comparison with respect to BIDAY queries response time

These normalized calculations provide the basis for comparing per-core efficiency and workload capacity across processor generations, which is reflected in the performance results discussed in the following sections.

CPU utilization and I/O throughput

This section summarizes CPU utilization and I/O characteristics observed during execution of the BIDAY workload using the traditional row format and provides supporting context for the capacity results presented in the section Capacity comparison of Power11 with Power9 running BIDAY queries.

The BIDAY queries running on row-format data are fully CPU intensive and moderately I/O intensive. An average CPU utilization of 99% to 100% was observed even for a single-user scenario on both Power11 and Power9 LPARs. While executing the BIDAY queries, the data throughput from the SAN volumes remained within the available bandwidth limits, and the I/O wait time on the host was minimal and within acceptable limits.

In addition to evaluating BIDAY query performance using the traditional row format, the evaluation also examines the impact of Oracle Database In-Memory on analytical workload performance.

Oracle Database In-Memory

This section describes the configuration and performance characteristics of the BIDAY workload when Oracle Database In-Memory is enabled, highlighting the impact of columnar storage on analytical query execution and per-core efficiency.

Oracle Database In-Memory is a powerful feature (licensed separately with Enterprise Edition) designed to accelerate analytic queries by storing data in a compressed, columnar format directly in memory. Unlike traditional row-based storage, this dual-format architecture allows data to be simultaneously available in both row and column formats, optimizing performance for both analytical and transactional workloads.

The database initialization parameters for the In-Memory database are provided in the Appendix section. The fact table sales_fact in BIDAY is loaded into the In-Memory area using the following SQL statement:

```
alter table gosldw.sales_fact inmemory priority critical;
```

This uses the default data compression option MEMCOMPRESS FOR QUERY LOW and completes loading the entire table into the In-Memory area during database startup.

The database initialization parameters for the In-Memory database are provided in the Appendix section.

For the BIDAY queries running on the data loaded into the in-memory area, 12 Power9 cores and 8 Power11 cores were assigned to the LPARs, based on the rPerf equation. When running queries with In-Memory enabled, a reduced number of cores was sufficient on both the Power9 and Power11 LPARs, as additional cores were primarily needed for row-format queries in the 1-, 2-, 4-, and 8-user scenarios. Consequently, with the reduced number of cores, CPU utilization ranged from approximately 75% for 1 user to 100% for 8 users.

The following table summarizes the per-core capacity improvement observed when running the BIDAY workload with Oracle Database In-Memory enabled on Power11 systems compared to Power9 systems.

Capacity metrics: Per-core improvement in the execution of queries in DB In-Memory	BIDAY workload with number of concurrent users			
	1 user	**2 users**	**4 users**	**8 users**
Relative per core capacity improvement factor ratio with **8** Power11 cores over **12** Power9 cores	2.57x	2.64x	2.59x	2.54x
Relative per core capacity improvement factor ratio with **6** Power11 cores over **12** Power9 cores	2.26x	2.10x	2.04x	1.97x

Table 5: Capacity comparison of Power11 with Power9 for BIDAY on Oracle Database In-Memory

The comparison of capacity metrics between Power11 cores and Power9 cores with respect to query execution time is shown in the following table.

Capacity metrics: Percentage increase in the execution of queries in DB In-Memory	BIDAY workload with number of concurrent users			
	1 user	**2 users**	**4 users**	**8 users**
Relative capacity increase percentage with **8** Power11 cores over **12** Power9 cores	22.31%	24.33%	22.80%	21.18%
Relative capacity increase percentage with **6** Power11 cores over **12** Power9 cores	11.60%	4.72%	1.85%	-1.59%

Table 6: Capacity comparison of Power11and Power9 with respect to BIDAY queries response time on Oracle Database In-Memory

When the number of Power11 processors was reduced by half, a higher load was observed in the 8-user scenario. As a result, the relative capacity decreased with respect to the response time to approximately -1.59%, and CPU utilization reached 100%. In contrast, for the 1-user scenario, CPU cycles were utilized

up to 75% on the Power9 LPAR and 85% on the Power11 LPAR. For the 2-user scenario and above, CPU utilization reached 100% in Power11 LPAR.

Conclusion

Based on experimental evaluations using BIDAY queries and online transaction processing (OLTP) workloads on Power11 and Power9 servers, IBM Power11 cores demonstrate significant capacity improvements over previous-generation Power cores. For complex workloads, Power11 cores can handle equivalent tasks with **half the number of cores**, delivering substantially higher per-core capacity for both row-format and in-memory queries under moderate to high stress. The increased capacity and reduced core count of Power11 systems can **potentially lower Oracle Database licensing costs, reduce power consumption, modernize database infrastructure with minimal risk and operational disruption, and decrease associated operational expenses**.

Appendix

Oracle Database 19c parameters used for BIDAY workload:

```
audit_file_dest='/u01/app/oracle/admin/biday/adump'
audit_trail='db'
compatible='19.0.0'
control_files='+DATA/BIDAY/CONTROLFILE/current.273.1204730001'
db_block_size=8192
db_create_file_dest='+DATA'
db_domain='ibpm.ihost.com'
db_name='biday'
diagnostic_dest='/u01/app/oracle'
dispatchers='(PROTOCOL=TCP) (SERVICE=bidayXDB)'
enable_pluggable_database=true
local_listener='LISTENER_BIDAY'
nls_language='AMERICAN'
nls_territory='AMERICA'
open_cursors=400
pga_aggregate_target=67108864000
processes=10240
remote_login_passwordfile='EXCLUSIVE'
sga_max_size=268435456000
sga_target=268435456000
undo_tablespace='UNDOTBS1'
_always_vector_transformation=TRUE
_star_transformation_enabled='TRUE'
parallel_degree_policy='AUTO'
optimizer_dynamic_sampling=0
optimizer_adaptive_reporting_only=TRUE
parallel_degree_limit=256
parallel_max_servers=1600
parallel_min_servers=256
recyclebin='OFF'
resource_manager_plan=''
```

The Database In-Memory parameters for the Oracle Database 19c version 19.27.0.0:

```
audit_file_dest='/u01/app/oracle/admin/biday/adump'
audit_trail='db'
compatible='19.0.0'
control_files='+DATA/BIDAY/CONTROLFILE/current.273.1204730001'
db_block_size=8192
db_create_file_dest='+DATA'
db_domain='pbm.ihost.com'
db_name='biday'
diagnostic_dest='/u01/app/oracle'
dispatchers='(PROTOCOL=TCP) (SERVICE=bidayXDB)'
enable_pluggable_database=true
local_listener='LISTENER_BIDAY'
nls_language='AMERICAN'
nls_territory='AMERICA'
open_cursors=400
pga_aggregate_target=67108864000
processes=10240
remote_login_passwordfile='EXCLUSIVE'
undo_tablespace='UNDOTBS1'
_always_vector_transformation=TRUE
_star_transformation_enabled='TRUE'
parallel_degree_policy='AUTO'
optimizer_dynamic_sampling=0
optimizer_adaptive_reporting_only=TRUE
parallel_degree_limit=64
parallel_max_servers=1600
parallel_min_servers=64
recyclebin='OFF'
resource_manager_plan=''
optimizer_inmemory_aware=TRUE
inmemory_optimized_arithmetic=ENABLE
inmemory_max_populate_servers=128
inmemory_size=260g
sga_target=400G
sga_max_size=400G
```

Oracle 19c Database parameters used for Brokerage OTLP workload:

```
*.ash_size=536870912
*.audit_file_dest='/u01/base/oracle/admin/ecomm/adump'
*.audit_trail='db'
*.compatible='19.0.0'
*.control_files='+DATA/ECOMM/CONTROLFILE/current.257.1211460631'
*.db_block_size=8192
*.db_name='ecomm'
*.diagnostic_dest='/u01/base/oracle'
*.dispatchers='(PROTOCOL=TCP) (SERVICE=ecommXDB)'
*.local_listener='LISTENER_ECOMM'
*.nls_language='AMERICAN'
*.nls_territory='AMERICA'
*.open_cursors=300
*.parallel_max_servers=600
*.parallel_min_servers=8
*.pga_aggregate_target=30G
*.processes=5000
*.remote_login_passwordfile='EXCLUSIVE'
*.resource_manager_plan=''
*.sessions=5505
*.sga_target=450G
*.undo_tablespace='UNDOTBS1'
```

Get more information

For more information on using Oracle Database on IBM Systems, you can contact ibmoracle@us.ibm.com.

- IBM Power11 server E1180
- Oracle Database 19c
- IBM FlashSystem Storage 7300

About the authors

Wayne Martin is the IBM Systems Technology Solutions Manager responsible for the technological 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 the 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.

Bhargavaram Akula is a Technical Consultant with IBM India, Hyderabad. He collaborates with the specialists at the IBM Oracle International Competency Center based in Redwood Shores (California, US) working on Oracle product certifications on IBM Power running AIX. He has extensive experience with Oracle products.

Ravisankar Shanmugam is an Advisory Software Engineer with the IBM Power, Oracle ISV engineering team in the IBM Infrastructure group based in San Jose, California. He writes technical collateral and leads proof-of-concept and performance evaluation projects for IBM Power servers with Oracle products.

Acknowledgment

IBM acknowledges Oracle engineering staff who provided guidance on the development and review of the white paper.

© Copyright IBM Corporation 2026

IBM Corporation
One Orchard Road
Armonk, NY 10504

Produced in the
United States of America
April 2026

IBM, the IBM logo, ibm.com, IBM Power, Power, Power9, Power11, and AIX are trademarks or registered trademarks of International Business Machines Corporation, in the United States and/ or other countries. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on ibm.com/trademark.

Oracle and Oracle Database are registered trademarks of Oracle Corporation and/or its affiliates.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates. All client examples cited or described are presented as illustrations of the manner in which some clients have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions. Generally expected results cannot be provided as each client's results will depend entirely on the client's systems and services ordered.

It is the user's responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs. THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.