International Journal of Computational Research and Development (IJCRD)
Impact Factor: 5.015, ISSN (Online): 2456 - 3137
(www.dvpublication.com) Volume 8, Issue 2, 2023

DATABASE HARDWARE CONFIGURATION AND ENERGY EFFICIENCY IN GLOBAL TRANSACTION PROCESSING

Divya Vishwanath* & Paul Johnson Asamoah**
* Janardhan Rai Nagar Rajasthan Vidyapeeth, Udaipur, Rajasthan, India
** Power Life University College, Accra, Ghana, West Africa

Cite This Article: Divya Vishwanath & Paul Johnson Asamoah, "Database Hardware Configuration and Energy Efficiency in Global Transaction Processing", International Journal of Computational Research and Development, Volume 8, Issue 2, July - December, Page Number 103-112, 2023.

Abstract:

This analysis examines whether larger database hardware configurations deliver greater transaction processing energy efficiency, addressing the distinction between computing capacity and energy productivity. Public Transaction Processing Performance Council records provide seven distinct configurations reported between 2010 and 2012. Descriptive comparisons evaluate processor cores, installed memory, and storage spindles against transactions per joule, calculated from matched throughput and full load power measurements. Energy efficiency ranges from 0.1489 to 1.4710 transactions per joule, a difference of 9.88 times. Equal processor core counts accompany an efficiency difference of 8.59 times, while larger memory capacities and spindle counts do not consistently indicate greater efficiency. The configuration with the highest throughput ranks fourth in energy efficiency. Comparisons distinguish transaction output from the power required to sustain it across complete configurations. Sensitivity checks preserve the efficiency ranking when the rounded official energy metric replaces the calculated measure. The contribution qualifies capacity centred interpretations through a traceable comparison of resource provision and achieved productivity. Procurement and capacity planning should therefore compare matched workload performance and power within consistent system boundaries.

Key Words: Benchmarking; Database Configuration; Descriptive Analysis; Energy Efficiency; Transaction Processing.

1. Introduction:

Database hardware investment can increase transaction capacity without producing a corresponding improvement in energy efficiency. This distinction matters when operators must meet processing requirements while managing electricity demand. Processor cores, memory, and storage expand the resources available to execute transactions, but their contribution depends on the throughput achieved relative to the power consumed by the complete configuration. The substantive question is therefore whether larger hardware configurations consistently deliver more transactions per joule. Answering this question requires separating hardware capacity, processing performance, and energy productivity.

The wider computing literature establishes why service growth and energy consumption require separate measurement. Masanet et al. 2020 show that improvements in infrastructure efficiency can moderate the relationship between expanding computing services and aggregate energy demand. Their analysis concerns global data centres rather than individual transaction processing configurations, but it identifies an important measurement principle: greater computational activity does not imply proportionately greater energy consumption. At the database level, this principle directs attention toward completed transactions relative to measured power rather than hardware size alone.

A resource allocation perspective explains why hardware expansion can produce competing outcomes. Additional processor cores can increase concurrent execution, provided memory and storage can sustain the additional activity. Greater memory capacity can reduce storage access when the execution workload uses that capacity effectively. Additional storage spindles can distribute input and output operations, while also adding equipment that consumes power. Beloglazov et al. 2012 develop resource allocation principles that connect computing demand with energy efficient provision. Applied to database configurations, this perspective implies that the usefulness of additional resources depends on their contribution to transaction output relative to their associated power requirements.

The closest database research provides evidence that these relationships depend on execution conditions. Zhou et al. 2020 demonstrate that selecting an appropriate number of processor cores can improve database energy efficiency by limiting memory swapping. Their results establish that processor capacity and supporting memory conditions interact operationally, rather than contributing independently through resource quantity alone. Mahajan et al. 2019 show that query optimization changes performance and energy consumption across relational and NoSQL databases. Together, these findings explain why the same hardware quantity can coexist with different energy outcomes. They also distinguish experimental interventions in database operation from comparisons among complete configurations whose components and software differ simultaneously.

Measurement boundaries are central to interpreting such comparisons. Orgerie et al. 2014 examine energy efficiency across computing and network resources, emphasizing the range of components and management choices involved in system energy use. For transaction processing, a throughput measure becomes

informative about energy productivity only when it is matched with a clearly defined power boundary. Processor core count, memory capacity, and spindle count describe available resources. Transactions per joule describes the output achieved relative to energy use. Treating these measures as interchangeable can obscure differences between configurations that provide greater capacity and those that use energy more efficiently.

The reviewed experimental evidence establishes that database operation and resource management can change energy efficiency. It does not determine whether processor cores, installed memory, and spindle counts consistently order the complete historical benchmark configurations examined here. This analysis addresses that bounded empirical question through a consolidated comparison of public Transaction Processing Performance Council records. Its contribution lies in connecting three reported hardware dimensions with a common energy productivity measure while preserving configuration identity, source provenance, and the reported power boundary.

The dataset contains seven distinct configurations drawn from eight energy reporting submissions. The retained reports cover 2010 through 2012 within an eligibility period extending from 2010 through 2022. Each configuration contributes one observation containing processor core count, installed memory, storage spindle count, transaction throughput, and full load power. One submission repeating the analytical hardware and energy measurements of another was excluded. The resulting cross section supports descriptive comparisons across configurations rather than estimates of isolated component effects.

Energy efficiency is calculated as transaction throughput divided by full load power for the reported energy configuration. The numerator measures transactions per second, and the denominator measures joules per second, producing transactions per joule. This construction retains the available precision of the reported inputs and distinguishes calculated efficiency from the rounded official energy metric. Matching throughput and power within each submission provides a transparent basis for comparison without treating larger capacity as evidence of better energy performance.

This study aims to describe how transaction processing energy efficiency varies across processor capacity, memory capacity, and storage configuration in the observed benchmarks. The central research question asks whether these hardware dimensions provide a consistent ordering of transactions per joule. The specific objectives are to compare efficiency across processor core counts, examine its variation across installed memory capacities, and assess its distribution across storage spindle counts. These objectives use the measured dimensions directly and retain complete configurations as the unit of interpretation.

The findings show that none of the three hardware dimensions independently provides a consistent efficiency ordering. Equal core counts accompany an efficiency difference exceeding eightfold, and the configuration with the highest throughput does not achieve the highest transactions per joule. These observations qualify the use of hardware capacity as a proxy for energy productivity. Their practical relevance is to comparative evaluation: operators should examine matched throughput and power measurements when screening candidate configurations. The transferable contribution is this evaluation principle, while the observed rankings remain specific to the represented systems, benchmark conditions, and report years.

2. Theory Conceptual Framework and Research Questions Development:

This analysis adopts a resource allocation perspective to examine database hardware configuration and transaction processing energy efficiency. Processing cores, memory, and storage operate as interdependent resources. Increasing one resource can improve transaction throughput, but the benefit depends on whether other components can support the additional work. Beloglazov et al. 2012 establish that energy efficient computing requires matching resource provision to computational demand. This provides the theoretical basis for examining hardware balance rather than treating greater capacity as an automatic source of efficiency.

The focal outcome is transaction processing energy efficiency, measured as reported transaction throughput divided by full load power consumption. Because watts measure joules per second, this ratio expresses completed transactions per joule. Hardware expansion improves this outcome only when throughput increases proportionately more than power consumption. Masanet et al. 2020 demonstrate at the broader infrastructure level that growth in computing services need not produce equivalent growth in energy consumption. That finding motivates the distinction between computing capacity and energy efficiency, without establishing the relationship for individual database configurations.

We operationalize database hardware configuration through three observed dimensions: processor core count, installed memory in gigabytes, and storage spindle count. These measures describe distinct physical resources and remain separate in the analysis. Energy efficiency covers the complete reported energy configuration, including database servers, storage, application servers, and miscellaneous equipment. The conceptual relationship therefore connects database hardware characteristics with efficiency at the reported system boundary.

Processor capacity represents the number of database server cores available for concurrent execution. Additional cores can increase completed transactions when sufficient parallel work exists. However, coordination overhead and competition for memory can restrict the throughput gained from additional processing resources. Zhou et al. 2020 show that selecting an appropriate number of cores can improve database

energy efficiency by avoiding memory swapping. Their experimental evidence supports a conditional mechanism: additional cores are useful when the supporting resources can sustain their activity. It does not establish that configurations with more cores will always complete more transactions per joule.

We examine processor capacity by comparing observed energy efficiency across the recorded core counts. The question concerns differences among complete configurations, so comparisons retain the reported system identity.

RQ1: How does transaction processing energy efficiency vary across processor core counts in the observed benchmark configurations?

Memory capacity represents the installed space available to hold database pages and execution structures. Greater capacity can reduce storage access when frequently used data remain in memory. Its contribution depends on the workload and software processes that determine how effectively that capacity is used. Mahajan et al. 2019 demonstrate that query optimization changes the performance and energy consumption of relational and NoSQL databases. Their evidence establishes that software choices can alter the energy consequences of a given hardware configuration. Installed memory therefore provides a measurable capacity indicator, while the theoretical benefit remains conditional on how the database uses it.

We compare energy efficiency across the observed memory capacities without treating memory size as a measure of cache effectiveness or memory utilization. The resulting question addresses a relationship that the recorded values can resolve.

RQ2: How does transaction processing energy efficiency vary across installed memory capacities in the observed benchmark configurations?

Storage configuration represents the reported number of storage spindles. Additional spindles can distribute input and output activity across devices, potentially increasing transaction throughput. They also add equipment that consumes power. The efficiency outcome depends on whether the throughput benefit exceeds the associated increase in power consumption. Orgerie et al. 2014 identify resource allocation, scheduling, and component operation as interacting determinants of computing energy efficiency. This systems perspective supports examining spindle count alongside the other hardware dimensions, while allowing efficiency to increase, decrease, or remain similar across configurations.

We use spindle count as the recorded storage characteristic. It does not represent storage capacity, device performance, or the number of every storage component.

RQ3: How does transaction processing energy efficiency vary across storage spindle counts in the observed benchmark configurations?

The conceptual framework connects processor capacity, memory capacity, and storage configuration with observed transaction processing energy efficiency. These connections represent descriptive associations among measured characteristics. We resolve the research questions through configuration level comparisons of core counts, memory capacities, spindle counts, and transactions per joule. The seven distinct configurations provide observations from 2010 through 2012 within the eligibility period of 2010 through 2022. Each configuration remains one observation, and the repeated hardware and energy record is excluded from the analytical sample.

3. Data:

This analysis documents the provenance, selection, and measurement of seven database benchmark configurations. The consolidated records support descriptive comparisons of processor capacity, memory capacity, storage configuration, and transaction processing energy efficiency.

3.1 Data Sources, Population, and Sample Construction:

The data originate from the Transaction Processing Performance Council public TPC E results register. The analytical workbook consolidates eight historical submissions containing hardware characteristics, transaction throughput, and reported energy measurements. Individual result identifiers provide the access route and link each extracted record to its source page. The workbook preserves source URLs and retrieval metadata separately from historical report dates. It does not identify a common benchmark specification revision for all submissions, so comparability rests on the reported fields and benchmark family.

The retained records comprise two HP configurations and five Fujitsu configurations. The HP reports document the ProLiant DL580 G7 and ProLiant DL585 G7, as reported by the Transaction Processing Performance Council, 2010a and 2010b. The Fujitsu reports cover the PRIMEQUEST 1800E2, PRIMERGY RX300 S6, and PRIMERGY RX900 S2 configurations, as reported by the Transaction Processing Performance Council, 2011a, 2011b, and 2011c. Two additional reports cover the PRIMERGY RX300 S7 and PRIMERGY RX500 S7, as reported by the Transaction Processing Performance Council, 2012a and 2012b. These institutional records supply the observations rather than estimates extracted from research figures.

The unit of analysis is a distinct hardware configuration with reported transaction throughput and energy consumption. The empirical setting is transaction processing under the TPC E benchmark. The target population comprises database configurations operating transaction workloads, while the accessible sampling frame comprises the eight energy reporting submissions identified in the public register within the eligibility

period of 2010 through 2022. The frame is international in vendor coverage but does not constitute a geographical sample of database installations.

The retained observations extend from 21 June 2010 through 5 November 2012. Two configurations were reported in 2010, three in 2011, and two in 2012. The data form a configuration cross section with irregular report dates. Each configuration contributes one observation. Hardware capacities describe the tested configuration, while throughput and power describe benchmark performance measurements. Report dates identify submissions rather than annual measurement intervals.

Selection required a report within the eligibility period and available values for processor cores, memory, storage spindles, throughput, and full load power. All eight extracted submissions met these field requirements. One submission repeated the hardware characteristics, throughput, and energy values of a later submission. Removing that repeated configuration produced seven analytical observations.

Table 1: Sample construction

Stage	Units	Observations	Reason or Rule
Initial accessible frame	8 submissions	8	Energy reporting submissions identified within the eligibility period
Repeated configuration excluded	1 submission	1	Result 111050601 repeats the analytical measurements reported in result 111060601
Consolidated analytical sample	7 configurations	7	Retain the later report for the repeated configuration
Missing analytical inputs excluded	0	0	All required measurements are populated
Final descriptive sample	7 configurations	7	Same observations used for all three hardware dimensions

The workbook records the eligibility and consolidation rules but contains no dated preregistration. We therefore describe these as implemented construction decisions without assigning them an undocumented prespecified status.

3.2 Data Integration, Cleaning, and Quality Assurance:

Integration involved appending fields from individual reports into a common schema. The TPC result identifier serves as the record key. Hardware characteristics, throughput, and energy values were matched within the same submission. No external dataset merge, temporal interpolation, or frequency conversion was required. Duplicate assessment distinguished submission identity from analytical repetition. Results 111050601 and 111060601 have different identifiers and report dates, but identical processor cores, memory, spindle count, throughput, full load power, and idle power. They remain separate records in the extraction sheet. The analytical sheet retains result 111060601, the later submission, to avoid counting the repeated hardware and energy configuration twice. This exclusion concerns the measured configuration rather than every field in the underlying reports.

The seven retained observations have complete values for all three hardware dimensions and both inputs used to calculate energy efficiency. Missingness is therefore zero across these fields, configurations, and represented report years. No missing values were replaced, interpolated, or assigned zero. Years without observations were not converted into empty annual records. Measurement checks confirmed positive core counts, memory capacities, spindle counts, throughput, and power denominators. Idle power remains below full load power in every retained configuration. The energy efficiency formula reproduces the stored outcome values from the corresponding throughput and power inputs. These checks establish internal consistency within the consolidated workbook.

Full load power was taken from the Total REC row in the energy subsystem table. This preserves the available precision rather than substituting the rounded power summary displayed elsewhere on the source page. The official watts per transaction per second metric remains a separate source field because rounding prevents its reciprocal from reproducing every calculated efficiency value exactly. All retained values remain in their reported units. No outlier threshold, trimming, or winsorization was applied. Consequently, large spindle counts and high-power measurements remain observable characteristics rather than automatic grounds for deletion. The source pages are historical records accessed through the public register; the workbook is not a reconstruction of information available to decision makers at each original report date.

The accessible frame depends on vendors reporting energy measurements. Selection into that frame can therefore differ from selection into the wider benchmark register. The resulting comparisons describe the retained configurations and do not represent the distribution of hardware or energy efficiency across global database installations.

3.3 Variable Measurement, Construction, and Validation:

Database hardware configuration is represented by three observed dimensions. Processor capacity is measured by database server core count. Memory capacity is measured by installed gigabytes. Storage

configuration is measured by reported spindle count. These physical quantities remain separate variables rather than components of a weighted index. Transaction processing energy efficiency is calculated using Equation 1.

$$Y_i = \frac{T_i}{P_i}$$

Here, Y_i denotes transactions per joule for configuration i, T_i denotes reported transaction throughput in transactions per second, and P_i denotes full load power in watts for the reported energy configuration. Dividing transactions per second by joules per second produces transactions per joule. Higher values indicate greater energy efficiency.

The outcome covers the reported energy configuration, including database servers, storage, application servers, and miscellaneous equipment. The hardware dimensions describe database system resources. Their relationship with the outcome therefore concerns complete tested configurations rather than isolated component efficiencies.

Table 2: Variable dictionary

Variable or Indicator	Analytical Role	Definition and Construction	Source And Identifier	Unit and Frequency	Transformation and Direction	Missing Value Rule	Coverage and Limitation
Processor core count	Hardware dimension	Reported database server total cores	TPC report, Total number of cores	Count per configuration	Original level; higher means more cores	No missing values	Seven observations; does not measure utilization
Installed memory	Hardware dimension	Reported database server memory size	TPC report, Memory Size	Gigabytes per configuration	Original level; higher means greater capacity	No missing values	Seven observations; does not measure memory access speed
Storage spindle count	Hardware dimension	Reported spindle count	TPC report, Spindle Count	Count per configuration	Original level; higher means more spindles	No missing values	Seven observations; does not measure storage capacity or every device
Transaction throughput	Outcome construction input	Reported TPC E throughput	TPC report, Throughput	Transactions per second during measurement	Original level; higher means greater throughput	No missing values	Seven observations; benchmark specific
Full load power	Outcome construction input	Total power for the reported energy configuration	TPC report, Total REC full load average watts	Watts during measurement	Original level; positive denominator	No missing values	Seven observations; includes supporting equipment
Transaction processing energy efficiency	Outcome	Throughput divided by full load power	Workbook formula using matched report inputs	Transactions per joule per configuration	Calculated ratio; higher means greater efficiency	All denominators positive	Seven observations; describes the reported system boundary
Official energy metric	Reference measurement	Published watts per transaction per second	TPC report, TPC Energy Metric	Watts per transaction per second	Original ratio; lower means greater efficiency	No missing values	Rounded source metric
Idle power	Contextual measurement	Total idle power for the reported energy configuration	TPC report, Total REC idle watts	Watts during idle measurement	Original level; higher means greater idle demand	No missing values	Seven observations; separate operating state
Result identifier and report date	Provenance	Submission identifier and historical report date	TPC report metadata	Identifier and calendar date	No transformation	No missing values	Link each observation to its source
Vendor and system	Configuration identity	Reported vendor and hardware model	TPC report metadata	Categories per configuration	No transformation	No missing values	Two vendors and seven configurations

Validation concerns the consistency of observed measurements and calculated ratios. The variables are physical quantities rather than latent constructs. Their interpretation follows the source field definitions and the documented energy boundary.

Table 3: Descriptive statistics

Variable	Valid Observations	Mean	Standard Deviation	Median	Minimum	Maximum
Processor cores	7	42.86	27.97	32.00	12.00	80.00
Memory in gigabytes	7	928.00	817.31	512.00	96.00	2,048.00
Storage spindles	7	264.57	253.64	120.00	68.00	752.00

Variable	Valid Observations	Mean	Standard Deviation	Median	Minimum	Maximum
Throughput in transactions per second	7	2,594.71	1,367.60	2,001.12	1,268.30	4,555.54
Full load power in watts	7	4,964.03	4,139.96	4,572.18	1,181.44	11,683.00
Energy efficiency in transactions per joule	7	0.8900	0.5424	0.9964	0.1489	1.4710

Standard deviations use the sample denominator. Statistics are calculated across configurations without throughput weighting. Each configuration therefore contributes equally to the reported summaries.

3.4 Analytical Framework, Identification, and Reproducibility:

The analysis is descriptive. Its object is the observed variation in transactions per joule across processor core counts, memory capacities, and storage spindle counts. Comparisons use the same seven configurations for each hardware dimension, preserving system identity and the reported measurement boundary.

The analytical framework treats hardware dimensions as concurrent characteristics of complete configurations. Differences between configurations can also reflect processor architecture, database software, supporting equipment, and benchmark implementation. Accordingly, observed differences describe configuration outcomes rather than separate effects attributable to individual hardware resources.

Reproducibility rests on the workbook structure. Extracted records preserve all eight submissions, selection labels identify the repeated configuration, source URLs retain report provenance, and the analytical sheet contains seven retained observations. The variable dictionary records definitions and units. Excel formulas expose the calculation of transactions per joule from the original throughput and full load power inputs.

All analytical measurements refer to the configuration documented in the same submission. No later hardware measurement is assigned to an earlier record. The shared sample, consistent units, and explicit denominator align the three hardware dimensions with the energy efficiency outcome. Interpretation remains bounded by the observed configurations, report years, and reported energy configuration.

4. Methods:

4.1 Research Design and Comparative Framework:

This analysis uses a descriptive cross sectional design to examine how transaction processing energy efficiency varies across complete database configurations. The analytical unit is a distinct benchmark configuration. The sample contains seven configurations reported between 2010 and 2012, following the selection process documented in Table 1.

The research questions concern observed differences across processor core counts, installed memory capacities, and storage spindle counts. Useful variation comes from differences in these characteristics across configurations. Comparisons retain system identity because each observation combines hardware, software, and supporting equipment within a reported benchmark environment.

The analytical target is the distribution and ordering of transactions per joule across the observed configurations. Comparisons describe complete systems rather than isolated component effects. Shared core counts or memory capacities provide informative contrasts, but they do not make configurations equivalent in their remaining characteristics. Interpretation therefore remains specific to the measured configurations and benchmark conditions.

4.2 Measurement and Descriptive Comparisons:

Energy efficiency is calculated using Equation 1, which divides transaction throughput by full load power for the same reported configuration. Throughput is measured in transactions per second and power in watts, producing transactions per joule. Higher values indicate greater transaction output per unit of energy. Calculations use the precision available in the source inputs, with rounding applied only to displayed results.

The three hardware dimensions retain their original units: processor cores, gigabytes of installed memory, and storage spindles. Definitions and measurement boundaries follow Table 2. These dimensions remain separate observed characteristics rather than components of a composite score.

Descriptive summaries report valid counts, arithmetic means, sample standard deviations, medians, minima, and maxima. The sample standard deviation uses the number of observations minus one as its denominator. Each configuration receives equal weight, so the summaries describe configurations rather than the distribution of transactions processed across systems.

Comparisons proceed separately for each hardware dimension. Configurations sharing the same core count or memory capacity are examined together, and their mean efficiency is calculated where multiple observations exist. Storage comparisons use the individual configurations because each retained spindle count is distinct. Throughput rankings are also compared with efficiency rankings to distinguish processing performance from energy productivity.

Practical magnitudes are expressed as ratios and percentage differences. A ratio divides the higher compared efficiency by the lower value. A percentage difference divides the difference between two efficiencies

by the stated reference configuration and multiplies the result by 100. Calculations use unrounded values. Tables 3 and 4 present the descriptive summaries and configuration comparisons.

4.3 Processing and Measurement Checks:

Processing begins with the eight extracted submissions and their source identifiers. The repeated analytical configuration is removed according to the rule recorded in Table 1, leaving seven observations. Throughput and power remain matched within each submission throughout outcome construction.

Completeness checks cover the three hardware dimensions and both inputs required by Equation 1. All retained observations contain these measurements. No imputation, interpolation, or replacement with zero is applied. Valid extreme observations remain in the sample, and no trimming or winsorization changes their contribution.

Internal consistency checks confirm positive hardware counts, memory capacities, transaction throughput, and full load power. Every outcome denominator is positive, and reported idle power is below full load power for each retained configuration. Recalculation from throughput and power verifies the stored efficiency values.

Measurement checks distinguish full load power from the rounded official energy metric. The principal outcome uses the Total REC full load power measurement, preserving its available precision. The official metric remains a separate reference field. These procedures assess arithmetic consistency and field alignment within the consolidated records.

4.4 Sensitivity Analysis and Reproducibility:

Two exploratory sensitivity checks address specific construction choices. The first restores the excluded repeated record and recalculates mean efficiency. This assesses how counting the repeated configuration changes the aggregate summary. The principal comparisons continue to use the seven distinct configurations.

The second replaces calculated efficiency with the reciprocal of the rounded official watts per transaction per second metric. It compares the resulting configuration ranking and percentage deviations against the principal measure. This check assesses sensitivity to source rounding while retaining the same seven observations.

Excel formulas implement the principal efficiency calculation. Python with NumPy supports descriptive summaries, grouped comparisons, rankings, and sensitivity calculations. These operations are deterministic and require no random initialization.

The workbook preserves extracted submissions, selection labels, source URLs, retained observations, variable definitions, and outcome formulas. Together, these materials allow the analytical sample and energy efficiency calculations to be reconstructed. The method answers the research questions through transparent comparisons of observed configurations, with conclusions bounded by the represented systems and measurement conditions.

5. Findings:

This analysis examines how transaction processing energy efficiency varies across processor capacity, memory capacity, and storage configuration. The findings resolve three descriptive research questions using seven distinct benchmark configurations. Configuration comparisons and measurement sensitivity checks establish the scope of the empirical conclusions.

The analytical sample contains two HP configurations and five Fujitsu configurations reported between 2010 and 2012. One submission with repeated hardware and energy measurements was excluded from the eight extracted records. All seven retained configurations have complete measurements for processor cores, memory, storage spindles, transaction throughput, and full load power. Energy efficiency averages 0.8900 transactions per joule, with a median of 0.9964 and a sample standard deviation of 0.5424. The highest observed efficiency is 9.88 times the lowest, indicating substantial variation among configurations operating within the same benchmark family.

Table 4: Hardware characteristics and transaction processing energy efficiency

Configuration	Processor Cores	Memory in Gigabytes	Storage Spindles	Transactions Per Second	Full Load Power in Watts	Transactions Per Joule
PRIMERGY RX500 S7	32	512	96	2,651.27	1,802.33	1.4710
PRIMERGY RX300 S7	16	512	68	1,871.81	1,288.82	1.4523
PRIMERGY RX300 S6	12	96	120	1,268.30	1,181.44	1.0735
PRIMERGY RX900 S2	80	2,048	384	4,555.54	4,572.18	0.9964

Configuration	Processor Cores	Memory in Gigabytes	Storage Spindles	Transactions Per Second	Full Load Power in Watts	Transactions Per Joule
PRIMEQUEST 1800E2	80	2,048	360	4,414.79	4,817.42	0.9164
HP ProLiant DL580 G7	32	1,024	752	2,001.12	11,683.00	0.1713
HP ProLiant DL585 G7	48	256	72	1,400.14	9,403.00	0.1489

For RQ1, processor core count does not order the configurations consistently by energy efficiency. The 16 core configuration achieves 1.4523 transactions per joule, exceeding both 80 core configurations, which achieve 0.9964 and 0.9164. The 48 core configuration records the lowest efficiency at 0.1489. More decisively, the two 32 core configurations differ by a factor of 8.59: the PRIMERGY RX500 S7 achieves 1.4710 transactions per joule, compared with 0.1713 for the HP ProLiant DL580 G7. Equal core counts therefore coexist with substantially different system outcomes. This pattern is consistent with the conditional relationship between processor resources and database energy efficiency examined experimentally by Zhou et al. 2020, although the present comparisons do not isolate processor effects.

For RQ2, greater installed memory does not correspond to a consistent increase in transactions per joule. The two configurations with 512 gigabytes achieve the highest efficiencies, averaging 1.4617 transactions per joule. Their individual values differ by only 1.29 percent, despite different processor and spindle counts. By comparison, the two configurations with 2,048 gigabytes average 0.9564 transactions per joule, while the configuration with 1,024 gigabytes achieves 0.1713. The 96-gigabyte configuration reaches 1.0735, exceeding all three configurations with memory capacities above 512 gigabytes. These comparisons establish that installed capacity alone does not reproduce the observed efficiency ranking. Mahajan et al. 2019 provide relevant experimental evidence that database optimization changes performance and energy consumption, supporting caution when interpreting memory capacity independently of the wider configuration.

For RQ3, spindle count also fails to produce a consistent efficiency ordering. The configurations with 68 and 96 spindles achieve 1.4523 and 1.4710 transactions per joule, respectively. The configuration with 72 spindles achieves only 0.1489, showing that similar spindle counts can accompany sharply different outcomes. At the upper end, the configuration with 752 spindles achieves 0.1713 transactions per joule. The two 80 core configurations, each containing 2,048 gigabytes of memory, provide a further descriptive contrast. The configuration with 384 spindles achieves 8.72 percent greater efficiency than the configuration with 360 spindles. Because these remain different complete systems, this comparison does not attribute the difference to storage. It reinforces the systems perspective developed by Orgerie et al. 2014, in which energy outcomes depend on interacting resources and operating arrangements.

Throughput and energy efficiency rankings also diverge. The PRIMERGY RX900 S2 records the highest throughput at 4,555.54 transactions per second but ranks fourth in transactions per joule. The PRIMERGY RX500 S7 ranks third in throughput and first in energy efficiency. The distinction arises because the efficiency measure evaluates completed transactions relative to total reported power demand. These observations support evaluating resource configurations against the intended performance and energy objective, consistent with the resource allocation perspective of Beloglazov et al. 2012. They do not establish an optimal hardware combination beyond the seven observed configurations.

Two exploratory sensitivity checks assess whether construction choices materially change the descriptive conclusions. Restoring the excluded repeated record increases mean efficiency from 0.8900 to 0.9033 transactions per joule, a change of 1.49 percent. The highest and lowest efficiency values remain unchanged. Separately, calculating efficiency as the reciprocal of the rounded official energy metric changes individual values by at most 0.37 percent and preserves the configuration ranking. Thus, the observed ordering persists under both checks, while the sample mean remains sensitive to counting a repeated configuration.

The three research questions yield a consistent descriptive answer: processor cores, installed memory, and spindle count each distinguish hardware capacity, but none independently provides a consistent ordering of transaction processing energy efficiency. The measured contribution is the comparison of complete configurations using a common throughput relative to power measure. These configuration results should remain distinct from aggregate infrastructure trends, such as those examined by Masanet et al. 2020, because seven benchmark configurations cannot establish changes in global energy demand.

6. Discussion:

Hardware capacity does not provide a reliable ordering of transaction processing energy efficiency across the observed configurations. Table 4 shows that equal processor core counts coexist with an efficiency difference exceeding eightfold, while greater memory and spindle counts do not consistently accompany higher transactions per joule. Equation 1 makes the distinction explicit: efficiency depends on throughput relative to

total reported power, rather than capacity alone. This analysis contributes a traceable comparison that separates hardware size, transaction performance, and energy productivity within a common benchmark family. The findings give configuration level substance to the resource allocation perspective of Beloglazov et al. 2012 by showing why resource quantities require interpretation alongside their measured system outcomes. The contribution is this empirical distinction, rather than evidence that any individual component determines efficiency.

The divergence between throughput and efficiency rankings changes how performance gains should be interpreted. The configuration delivering the highest throughput ranks fourth in energy efficiency, whereas the efficiency leader delivers less throughput with a smaller reported power requirement. Under Equation 1, greater throughput improves efficiency only when its proportional increase exceeds the associated increase in power. Table 4 therefore identifies different balances between transaction output and power demand without assigning those balances to an unmeasured process. Zhou et al. 2020 experimentally demonstrate that processor management and memory constraints can alter database energy efficiency. The present evidence complements that intervention-based perspective by showing that published configuration outcomes cannot be interpreted from core counts alone. Its contribution is to distinguish an observed system balance from the mechanisms that might produce it.

The hardware comparisons also refine the meaning of capacity within the conceptual framework. The two configurations containing 512 gigabytes of memory achieve similar efficiencies despite different processor and spindle counts, while the configurations containing the largest memory capacities perform less efficiently. This pattern makes installed memory an incomplete descriptor of energy performance rather than a direct measure of productive memory use. Likewise, similar spindle counts accompany sharply different outcomes in Table 4. Mahajan et al. 2019 show that database query optimization changes both performance and energy consumption, illustrating why hardware quantities cannot represent the full execution environment. The theoretical implication is that processor, memory, and storage dimensions describe available resources, while transactions per joule describes the outcome of their operation within a complete configuration.

Measurement boundaries and record selection materially shape the interpretation of benchmark evidence. Equation 1 uses power for the complete reported energy configuration, including supporting equipment, while the hardware dimensions describe selected database resources. The resulting efficiency measure therefore supports comparison of complete configurations rather than attribution of energy use to individual components. The exploratory checks further distinguish measurement precision from sample construction: using the rounded official energy metric preserves the efficiency ranking, while restoring the repeated record changes the sample mean. These results make provenance and denominator definition part of the substantive contribution. Orgerie et al. 2014 examine energy efficiency across interacting computing and network resources; the present analysis applies that systems perspective through an explicit measurement boundary and an auditable treatment of repeated observations.

The international relevance lies in a transferable evaluation principle: procurement and capacity planning should compare transaction throughput with power measured over a consistent system boundary. Table 4 demonstrates that selecting the largest observed processor, memory, or storage configuration would not identify the most energy efficient configuration in this sample. Decision makers should therefore request matched throughput and power measurements for candidate systems under relevant workloads. Masanet et al. 2020 distinguish growth in computing services from growth in aggregate energy use; this analysis addresses the corresponding distinction at the benchmark configuration level without extrapolating to national or global demand. Together, Table 4 and Equation 1 redirect attention from resource expansion toward measured energy productivity. The evidence supports comparative screening of the observed systems, while broader evaluation requires configurations and workloads beyond those represented here.

7. Conclusion and Implications:

Expanding digital infrastructure does not ensure that each unit of energy delivers more useful work. This analysis shows that larger resource configurations do not consistently achieve greater energy productivity, and that superior processing performance can coexist with lower efficiency. The contribution is a traceable comparison that distinguishes resource provision from productive energy use across complete benchmark configurations. This evidence qualifies capacity centred interpretations of computing efficiency by demonstrating that equipment quantities cannot substitute for measured system outcomes. Its theoretical value lies in treating configuration characteristics and achieved productivity as distinct analytical objects, while retaining the operational boundary within which transactions and power are measured. For managers, the findings support evaluating candidate systems against matched workload performance and power requirements before committing resources. Procurement teams should require comparable measurement boundaries and distinguish throughput requirements from efficiency objectives. Benchmark governance can strengthen these decisions through explicit reporting of supporting equipment, calculation precision, and repeated submissions. Operationally, retaining source identifiers and reproducible calculations makes comparisons easier to audit and update. The wider relevance is a transferable evaluation principle for organizations managing digital service

demand under energy constraints. Applying that principle can inform resource allocation without assuming that equipment expansion necessarily improves environmental performance.

The seven configurations provide a foundation for broader comparative inquiry. Additional vendors, database platforms, and workload types would extend the range of represented systems. Repeated observations under controlled operating conditions could distinguish hardware differences from software and workload influences. Direct component measurements would clarify how supporting equipment contributes to total demand. Experiments that vary resources individually and jointly could test the mechanisms behind the observed configuration differences. Longitudinal benchmarks could assess whether efficiency rankings persist across operating loads and technology generations. These extensions would establish where the evaluation principle transfers and where deployment conditions require different interpretations of energy productivity.

References:

1. Beloglazov, A., Abawajy, J., &Buyya, R. (2012). Energy-aware resource allocation heuristics for efficient management of data centers for Cloud computing. Future Generation Computer Systems, 28(5), 755-768. https://doi.org/10.1016/j.future.2011.04.017
2. Mahajan, D., Blakeney, C., & Zong, Z. (2019). Improving the energy efficiency of relational and NoSQL databases via query optimizations. Sustainable Computing: Informatics and Systems, 22, 120-133. https://doi.org/10.1016/j.suscom.2019.01.017
3. Masanet, E., Shehabi, A., Lei, N., Smith, S., &Koomey, J. (2020). Recalibrating global data center energy-use estimates. Science, 367(6481), 984-986. https://doi.org/10.1126/science.aba3758
4. Orgerie, A. C., de Assuncao, M. D., & Lefèvre, L. (2014). A survey on techniques for improving the energy efficiency of large-scale distributed systems. ACM Computing Surveys, 46(4), Article 47. https://doi.org/10.1145/2532637
5. Transaction Processing Performance Council. (2010a). HP ProLiant DL580 G7: TPC-E result details, result 110062202. https://www.tpc.org/tpce/results/tpce_result_detail5.asp?id=110062202
6. Transaction Processing Performance Council. (2010b). HP ProLiant DL585 G7: TPC-E result details, result 110062103. https://www.tpc.org/tpce/results/tpce_result_detail5.asp?id=110062103
7. Transaction Processing Performance Council. (2011a). PRIMEQUEST 1800E2: TPC-E result details, result 111072701. https://www.tpc.org/tpce/results/tpce_result_detail5.asp?id=111072701
8. Transaction Processing Performance Council. (2011b). PRIMERGY RX300 S6 12x2.5: TPC-E result details, result 111021401. https://www.tpc.org/tpce/results/tpce_result_detail5.asp?id=111021401
9. Transaction Processing Performance Council. (2011c). PRIMERGY RX900 S2: TPC-E result details, result 111060601. https://www.tpc.org/tpce/results/tpce_result_detail5.asp?id=111060601
10. Transaction Processing Performance Council. (2012a). PRIMERGY RX300 S7: TPC-E result details, result 112070501. https://www.tpc.org/tpce/results/tpce_result_detail5.asp?id=112070501
11. Transaction Processing Performance Council. (2012b). PRIMERGY RX500 S7: TPC-E result details, result 112110501. https://www.tpc.org/tpce/results/tpce_result_detail5.asp?id=112110501
12. Zhou, Y., Taneja, S., Qin, X., Ku, W.-S., & Zhang, J. (2020). EDOM: Improving energy efficiency of database operations on multicore servers. Future Generation Computer Systems, 105, 1002-1015. https://doi.org/10.1016/j.future.2017.02.043