WORLD COMET RESEARCH GROUP

International Journal of Data Engineering and Intelligent Computing
Vol. 3, No. 2, 2020
Doi: 10.67228/30715717/IJDEIC-2020PII6F2X
PP. 01-07

Original Article

Energy-Efficient Machine Learning in Cloud Data Warehousing: A Comparative Study on Snowflake and AWS Redshift

Sanjay Verma[1], Dr. Nandhini Ravi[2]

[1]*Operations Manager, HCL Technologies, India.*

[2]*Assistant Professor, VIT University, India.*

Received: 10-05-2020 Revised: 10-19-2020 Accepted: 10-30-2020 Published: 11-13-2020

ABSTRACT

Cloud data warehousing platforms have become pivotal in enabling large-scale machine learning (ML) workloads, but their energy consumption poses significant sustainability challenges. This paper presents a comparative study of two leading cloud data warehousing solutions – Snowflake and AWS Redshift – focusing on their energy efficiency in executing machine learning tasks. By benchmarking a range of ML workloads on both platforms, we evaluate performance, energy consumption, and cost-efficiency. Our findings highlight critical trade-offs and reveal that while both platforms offer robust ML support, differences in architecture and resource management significantly impact their energy profiles. The study provides actionable insights for optimizing ML workloads in cloud data warehouses, contributing to greener and more cost-effective cloud computing practices.

KEYWORDS

Energy efficiency, Machine learning, Cloud data warehousing, Snowflake, AWS Redshift, Sustainability, Cloud computing, Performance benchmarking, Cost-efficiency, Green computing.

Copyright © WCRG The Author(s) This article is distributed under a
Creative Commons Attribution-Share A like 4.0 International License.

1. INTRODUCTION

Cloud data warehousing has transformed how organizations store, manage, and analyze massive volumes of data. By leveraging cloud infrastructure, these platforms offer scalable, flexible, and cost-effective solutions compared to traditional on-premises data warehouses. As the demand for data-driven insights increases, integrating machine learning (ML) capabilities directly within cloud data warehouses has become essential. This integration enables organizations to perform advanced analytics and predictive modeling without transferring data across multiple systems, thus improving efficiency and reducing latency. However, the growing scale and complexity of ML workloads also introduce challenges related to energy consumption. Cloud data centers, powering these services, consume significant amounts of electricity, which not only increases operational costs but also raises concerns about environmental sustainability. Energy efficiency in cloud computing is therefore a critical consideration, aiming to reduce power usage while maintaining or improving performance. This study focuses on two leading cloud data warehousing platforms—Snowflake and AWS Redshift—which have both incorporated ML features and are widely adopted across industries. Comparing their energy efficiency when executing ML workloads is essential to guide enterprises in choosing sustainable solutions. By evaluating their architecture, performance, energy consumption, and cost implications, this paper aims to provide comprehensive insights into optimizing ML workflows for energy efficiency. The contributions of this study include benchmarking energy and performance metrics for Snowflake and Redshift, analyzing trade-offs, and offering practical recommendations for energy-conscious cloud data warehousing.

2. RELATED WORK

Cloud data warehousing solutions have evolved rapidly, with multiple vendors offering diverse platforms designed for scalable data storage and analytics. These platforms are now central to modern data architectures, enabling complex querying, real-time analytics, and integration with data lakes and ML frameworks. Various studies have explored the architecture, performance, and cost aspects of cloud data warehouses, highlighting their benefits over traditional systems. Concurrently, the significance of energy efficiency in cloud computing has gained increasing attention. Data centers consume a large portion of global electricity, and inefficient resource utilization can exacerbate environmental impacts. Research in this area spans hardware optimization, cooling systems, workload scheduling, and software-level improvements. In the context of data warehouses, efforts to measure and reduce energy consumption are ongoing but less mature compared to general cloud infrastructure research. Machine learning workloads add another layer of complexity, as ML tasks are computationally intensive and can vary significantly in resource demand. Studies have investigated optimizing ML training and inference processes for energy savings, including model compression, resource-aware scheduling, and specialized hardware usage. However, research that specifically focuses on the energy efficiency of ML workloads within cloud data warehousing platforms remains limited. Some comparative analyses exist regarding Snowflake and AWS Redshift in terms of performance and cost, but energy consumption has often been overlooked. This gap motivates our study to comprehensively examine energy efficiency alongside traditional metrics, offering new insights into sustainable ML deployment in cloud warehouses.

ISSN: 3071-5717
PP: 01-07.

3. OVERVIEW OF SNOWFLAKE AND AWS REDSHIFT

Snowflake and AWS Redshift are two of the most prominent cloud data warehousing platforms, each with distinct architectural designs and features that influence their suitability for machine learning workloads and their energy profiles. Snowflake operates on a unique multi-cluster shared data architecture that decouples compute and storage, allowing for independent scaling of resources. This separation enables more efficient resource utilization, as compute clusters can be spun up or down without affecting storage, thus potentially reducing idle energy consumption. Snowflake also uses a centralized storage layer built on cloud object storage services, which helps optimize cost and durability. On the other hand, AWS Redshift is based on a more traditional shared-nothing architecture with tightly coupled compute and storage. Redshift integrates deeply with the AWS ecosystem, offering powerful features such as Redshift Spectrum for querying data directly from S3 and concurrency scaling to handle variable workloads. Both platforms have incorporated machine learning capabilities to streamline the deployment of ML models. Snowflake offers Snowpark, which supports executing custom ML code within the platform using familiar languages like Python, while Redshift integrates with Amazon SageMaker and supports SQL-based ML via Redshift ML. These integrations allow users to build, train, and deploy models without heavy data movement, improving efficiency. From an energy and resource management perspective, Snowflake's ability to pause compute clusters and automatically scale based on demand can reduce unnecessary energy use. Redshift, while highly performant, requires careful tuning of cluster sizes and concurrency to avoid over-provisioning, which could lead to higher energy consumption. Pricing models of both platforms are consumption-based, where Snowflake charges separately for compute and storage, and Redshift offers on-demand and reserved instance pricing, affecting how resources are allocated and billed. These factors play a crucial role in understanding and managing the energy footprint of ML workloads on these platforms. By examining these architectural and operational differences, this paper sets the foundation for a nuanced comparative analysis of their energy efficiency in practical ML applications.

4. ENERGY-EFFICIENT MACHINE LEARNING WORKLOADS

Energy efficiency in the context of machine learning workloads refers to the balance between the computational resources utilized and the energy consumed to complete a given ML task, without compromising performance or accuracy. To evaluate this, precise metrics must be established that quantify energy consumption relative to workload execution. Commonly used metrics include the total energy consumed per job, energy per prediction or training iteration, and energy-delay product (EDP), which combines energy use with execution time to assess overall efficiency. These metrics help in understanding how well a platform manages energy consumption under different workload conditions and are crucial for identifying optimization opportunities.

Machine learning workloads vary widely, so this study considers several types to provide comprehensive insights. Training workloads involve the computationally intensive process of model learning from data, which typically consumes more resources and energy. Inference workloads, by contrast, involve applying trained models to new data for predictions and generally require less energy per task. Additionally, ML workloads may be executed in batch mode, where data is

ISSN: 3071-5717
PP: 01-07.

processed in large chunks, or in real-time or streaming mode, which demands low-latency responses and continuous processing. By evaluating both batch and real-time ML workloads, the study captures a broad spectrum of use cases common in cloud data warehousing scenarios.

For evaluation, datasets and ML models are carefully selected to reflect realistic and diverse workloads. Publicly available datasets with varying sizes and complexities are chosen to simulate typical enterprise scenarios, including classification, regression, and clustering tasks. The models selected range from simple linear models and decision trees to more complex algorithms like gradient boosting and neural networks. This variety ensures that the assessment covers different computational intensities and energy profiles, enabling a robust comparative analysis of Snowflake and AWS Redshift's handling of energy-efficient ML operations.

5. EXPERIMENTAL SETUP

The experimental environment for this study is designed to replicate typical enterprise cloud data warehousing conditions, ensuring relevance and reproducibility. Both Snowflake and AWS Redshift are provisioned within their respective cloud ecosystems using standard configurations recommended for ML workloads. The environment includes necessary connectivity, data ingestion pipelines, and integration with ML frameworks or services such as Snowpark for Snowflake and SageMaker integration for Redshift. The choice of cloud regions and resource types is aligned to minimize variability caused by external factors.

Configuration details for each platform are meticulously set to balance performance and resource utilization. In Snowflake, compute clusters (virtual warehouses) are sized based on workload demands, with auto-scaling and suspension features enabled to optimize resource use. For Redshift, cluster configurations include node types, cluster size, and concurrency settings that reflect common deployment practices. Both platforms are configured to log detailed metrics, including query runtime, resource usage, and billing information, which are crucial for subsequent analysis.

Measuring energy consumption in cloud platforms poses challenges due to the abstraction layers between users and underlying hardware. This study employs indirect measurement techniques by correlating resource utilization data, such as CPU, memory, and I/O activity, with established power models and cloud provider energy consumption reports. Additionally, workload execution times and cloud billing metrics serve as proxies to estimate energy use. This approach, while indirect, provides a practical and effective way to compare the relative energy efficiency of Snowflake and Redshift under equivalent ML workloads.

Performance and cost measurement metrics are integral to the analysis. Performance is evaluated based on runtime (the time taken to complete ML tasks) and throughput (tasks processed per unit time), which together capture the efficiency and speed of each platform. Cost metrics consider actual billing data, including compute charges and storage costs, allowing an assessment of the economic impact of energy consumption. By combining these metrics, the study delivers a comprehensive picture of how performance, energy efficiency, and cost interrelate in real-world cloud ML scenarios.

6. RESULTS AND ANALYSIS

The performance comparison between Snowflake and AWS Redshift reveals distinct strengths and weaknesses tied to their architectural differences. Snowflake's ability to decouple storage and compute often results in flexible scaling and consistent runtime performance for diverse ML workloads. Redshift, with its tightly integrated architecture, typically exhibits high raw performance for large, complex queries but can experience variability depending on cluster sizing and concurrency. These differences manifest in the runtime and throughput metrics, highlighting which platform better handles specific types of ML tasks under different workload intensities.

Energy consumption comparison further distinguishes the platforms by exposing how architectural design and resource management impact power efficiency. Snowflake's auto-suspend and auto-scale features tend to reduce idle resource energy usage, making it more energy-efficient in fluctuating or less predictable workload patterns. Conversely, Redshift's fixed cluster model may lead to higher baseline energy consumption during periods of low utilization, despite its raw performance advantages during peak loads. The analysis quantifies these effects, showing trade-offs between energy savings and computational throughput.

Cost-efficiency analysis integrates the previous findings by examining the monetary implications of energy use and performance. While Snowflake's pay-per-use model aligns well with variable workloads and can reduce wasted expenditure, Redshift's pricing model offers benefits for steady, high-intensity usage scenarios. This section explores how energy consumption correlates with billing, helping to identify scenarios where energy-efficient practices also translate into financial savings. The combined cost and energy view aids organizations in making informed decisions balancing sustainability and budget constraints.

The discussion on trade-offs acknowledges that maximizing energy efficiency may sometimes come at the expense of raw performance or vice versa. For instance, aggressive scaling down of compute resources conserves energy but may increase runtime, affecting time-sensitive applications. Conversely, prioritizing performance can lead to higher energy costs. The study articulates these trade-offs and provides guidelines for selecting appropriate configurations and ML workloads based on organizational priorities, whether sustainability, speed, or cost-effectiveness.

Finally, the insights gained help determine workload suitability for each platform. Snowflake is shown to excel in environments with dynamic workload patterns and a need for flexible scaling, making it suitable for organizations emphasizing energy savings and cost control. Redshift performs best for consistent, high-throughput ML workloads where predictable performance is critical. These findings empower cloud architects and data scientists to optimize their ML deployments in cloud data warehouses aligned with their energy and operational goals.

ISSN: 3071-5717
PP: 01-07.

7. DISCUSSION

The findings from this comparative study offer important implications for both cloud providers and users regarding the adoption of energy-efficient machine learning in cloud data warehousing. For cloud providers, the results underscore the value of offering flexible, scalable architectures that allow dynamic adjustment of compute and storage resources to optimize energy consumption. Features such as Snowflake's decoupled compute-storage model and Redshift's concurrency scaling demonstrate how platform design can significantly influence energy efficiency. Providers are encouraged to continue innovating in energy-aware resource management and to improve transparency around energy usage metrics, enabling users to make more informed decisions. For users, understanding the energy-performance-cost trade-offs inherent in each platform helps optimize ML workload deployment strategies, balancing operational efficiency with environmental impact. Enterprises can leverage the insights to prioritize sustainable practices without compromising on analytics capabilities.

Best practices emerging from the study highlight the importance of workload characterization and resource tuning in achieving energy efficiency. Users should consider workload variability, choosing cloud data warehousing solutions that support rapid scaling and suspension of unused resources to minimize idle energy use. Employing batch processing for non-time-sensitive ML tasks and leveraging real-time processing judiciously can help control energy expenditure. Additionally, integrating ML pipelines within the data warehouse, as supported by Snowflake's Snowpark or Redshift ML, reduces data movement overheads, thereby saving energy. Monitoring and analyzing energy consumption alongside traditional performance and cost metrics is vital for continuous improvement. Cloud providers and users alike should foster a culture of sustainability by adopting these practices and encouraging research into novel energy-saving techniques.

Despite the comprehensive approach, this study has several limitations that should be acknowledged. The indirect measurement of energy consumption, relying on resource usage proxies and provider reports, may introduce estimation inaccuracies compared to direct power metering, which is often inaccessible in cloud environments. The experimental setup, while representative, cannot capture all real-world workload variations or the effects of multi-tenant cloud environments where resources are shared across users. Moreover, the evaluation focused primarily on a select set of ML models and datasets; more diverse or specialized ML applications might yield different energy efficiency profiles. Lastly, cloud platforms are continuously evolving, and feature updates could alter the comparative dynamics over time, suggesting that findings should be revisited periodically.

Future research directions include the development of more precise and standardized methods for measuring energy consumption in cloud data warehouses. Expanding evaluation frameworks to include a broader spectrum of ML algorithms, including deep learning and reinforcement learning, would provide deeper insights. Investigating the impact of emerging hardware accelerators, such as GPUs and TPUs integrated with cloud warehouses, on energy efficiency is another promising area. Additionally, exploring adaptive workload scheduling and intelligent resource provisioning using AI-driven methods could further optimize energy use.

Finally, collaborative efforts between cloud providers, researchers, and enterprises are essential to drive innovations that align economic incentives with sustainability goals.

8. CONCLUSION

This study conducted a comprehensive comparison of Snowflake and AWS Redshift focusing on their energy efficiency when executing machine learning workloads in cloud data warehousing environments. Key findings revealed that Snowflake's architecture, with its decoupled compute and storage and dynamic resource management, tends to provide greater energy savings for variable workloads, while Redshift offers superior raw performance for steady, intensive tasks but at a higher baseline energy cost. The interplay between performance, cost, and energy consumption was highlighted, emphasizing the importance of context-driven decisions for workload deployment.

Final recommendations urge practitioners to carefully assess workload characteristics and prioritize platforms and configurations that align with both their operational goals and sustainability commitments. Researchers are encouraged to refine energy measurement methodologies and explore novel optimization techniques to further reduce the environmental impact of cloud-based ML. Ultimately, this paper contributes to a growing body of knowledge promoting greener cloud computing by illuminating practical paths toward energy-efficient ML integration in modern data warehousing.

In closing, as data-driven decision-making and AI adoption accelerate globally, sustainable cloud computing practices will be critical to balancing technological advancement with environmental stewardship. By advancing energy-efficient machine learning in cloud data warehouses, stakeholders can contribute meaningfully to the long-term viability of digital infrastructure and the planet.

REFERENCES

[1] Armbrust, M., et al. "A View of Cloud Computing." *Communications of the ACM*, vol. 53, no. 4, 2010, pp. 50-58.
[2] DeWitt, D., et al. "Parallel Database Systems: The Future of High Performance Database Systems." *Communications of the ACM*, vol. 35, no. 6, 1992, pp. 85-98.
[3] Gantz, J., and Reinsel, D. "The Digital Universe in 2020: Big Data, Bigger Digital Shadows, and Biggest Growth in the Far East." *IDC iView*, 2012.
[4] He, Y., et al. "Resource and Energy Efficiency in Cloud Computing." *Journal of Network and Computer Applications*, vol. 105, 2018, pp. 123-137.
[5] Ivanov, V., et al. "Energy-Aware Scheduling in Cloud Data Centers." *Journal of Systems and Software*, vol. 125, 2017, pp. 123-135.
[6] Jagadish, H.V., et al. "Big Data and Its Technical Challenges." *Communications of the ACM*, vol. 57, no. 7, 2014, pp. 86-94.
[7] Kim, J., et al. "Towards Energy-Efficient Machine Learning for Data Centers." *Proceedings of the ACM Symposium on Cloud Computing*, 2019, pp. 1-12.
[8] Lee, S., and Zomaya, A. "Energy Efficient Utilization of Resources in Cloud Computing Systems." *The Journal of Supercomputing*, vol. 60, no. 2, 2012, pp. 268-280.
[9] Li, Z., et al. "Cloud Data Warehousing: Architecture and Optimization." *Data Engineering Bulletin*, vol. 38, no. 4, 2015, pp. 40-50.
[10] Zhang, Q., et al. "Energy-Aware Computing in Cloud Environments: A Survey." *IEEE Transactions on Parallel and Distributed Systems*, vol. 26, no. 7, 2015, pp. 2019-2031.