

Benchmarking Cloud-Native ETL Tools Performance, Cost, and Scalability

¹Mr. Raghavendar Nellikondi,
Independent researcher, Sterling, VA, USA,
Email: raghav.nellikondi@gmail.com

²Mr. Reddaiah Kasturi
Independent researcher, Sterling VA, USA.
Email: Reddaiah.Kasturi@yahoo.com

Received: 20.10.2024

Revised: 16.11.2024

Accepted: 19.12.2024

Abstract

This work evaluates the performance, cost and their scalability for four popular cloud native ETL tools on public clouds. We were able to process ordered/equijoin/GroupBy datasets of size 10 GB up to 2 TB (both structured and semi-structured data) in a controlled experimental environment, using the same compute/network settings. Performance The performance of the solutions was measured in terms of end-to-end execution time, throughput (GB/min), and latency; the scalability was analyzed by observing the variation of performances when increasing from $1\times$ to $8\times$ parallel workloads. b, was performed using real cloud billing data in USD per terabyte processed (USD / TB). Our results show that serverless ETL tools outperform cluster-based ETL frameworks up to 32–45% for workloads smaller than 500 GB, and cluster-based frameworks are more stable in large-scale data processing with only 12% performance regression at peak time. Cost efficiency varied greatly between the solutions, serverless tools reduced operational costs by 28 to 40% for incomplete workloads and long running pipelines preferred managed clusters. The scalability tests detected nearly linear scaling efficacy (0.91–0.97) of the implemented tool in large concurrency conditions when tools were optimized. The results give developers and architects data as to which ETL tools fit the required workload size, cost limits, and scalability needs in cloud-native data ecosystems.

Keywords: Cloud-native ETL, Performance benchmarking, Cost efficiency, Scalability analysis, Data engineering.

1. Introduction

The growing popularity of data-driven applications has led to a dramatic increase in the demand for scalable, efficient and cost-effective data integration solutions. Today's businesses are processing a large amount of structured, semi-structured and unstructured data that comes from diverse sources including Internet-Of-Things (IoT) devices, transactional systems, social media applications and cloud-based applications. Extract, Transform and Load (ETL) operations are the key to preparing this data for analytics, machine learning or business intelligence pipelines [1]. Traditional ETL solutions that are not cloud-based can be challenged when facing these demands even with more effective tools, as these could be outdated or unsuitable to a Big Data scenario [2].

One promising solution has been the rise of cloud-native ETL tools which exploit inherent benefits of Cloud Computing like elastic provisioning, pay-as-you-go pricing models and managed infrastructure services [3]. It's built to fit well in the cloud system, scaling resources automatically and managing redundancy of hardware to make a Bash scripter proud. Since more and more enterprises are moving their data pipelines to the cloud, choosing an ETL tool is a strategic choice affecting both performance, operational cost and future scalability [4].

Even with the prevalence of cloud-native ETL tools, there are no comprehensive studies that provide quantitative comparisons of such solutions in a controlled and reproducible environment [5]. Current studies often emphasize either architectural betterment or qualitative comparisons, with little empirical explanation for efficiency of execution, scale-up behavior and trade-offs between cost-arrest [6]. Furthermore, benchmarks from vendors are often not transparent, preventing users from generalizing the results if their workloads have different patterns [7].

Performance is an important consideration when choosing an ETL tool, especially with analytics that are time sensitive or applications demanding real-time data feeds. Performance goals, such as time to run and throughput, can be used as measurable performance indicators of system efficiency given different amount of data or workloads [8]. Cost effectiveness is of equal importance because cloud pricing models pose new issues in relation to resource demand, idle costs, and workload burstiness [9]. Scalability Scalability also relates to the capability of a tool like ETLMate to handle increasing amounts of data without exponential cost or performance degradation.

10.48047/jocaaa.2024.33.08.394

In this setting, it is crucial to evaluate cloud native ETL tools with benchmark standard datasets and uniform infrastructure settings using objective performance metrics. This allows data engineers, system designers, and decision makers to align on how to wisely use the resources they have in context of the organizational need [10]. The present research gap is then addressed in this paper with a comparative (both systematic and quantitative) study of cloud-native ETL tools, considering the aspects of performance, cost and scalability. The results of this study intend to provide practical implications for both academia and industry, including easy development and deployment of cloud-based data integration pipelines.

2. Literature Review

Some previous works of cloud-based data integration focused on the drawbacks of legacy ETL systems for big and dynamic workloads. These researches emphasized difficulties like hard infrastructure, lack of scalability and high maintenance costs. The trend towards cloud-native ETL architectures were suggested as a way of responding to elasticity, fault-tolerance, and automatic resource management in data-intensive enterprise applications [11].

Later studies examined the concept models of cloud-native ETL tools at architectural level with an emphasis on microservice-oriented approaches, virtualization-container methods, and serverless processing behaviors. These results illustrated how decomposing compute and storage provides elasticity and fine-granularity. On the other hand they noted increased complexity in arrangement and unpredictability of performance under very high levels of concurrency [12].

Performance analysis of ETL tools in cloud based environments with synthetic and real datasets has been investigated in some studies. Indicators such as latency, throughput, pipeline delay, and performance were the most commonly used for measuring efficiency. Results showed that performance of different systems are quite diverse with respect to workload type, volume of data and resources setting which highlighted the importance of standard benchmarking [13].

Cost optimization becomes one of the major research hotspots after popularization of pay-as-you-go cloud pricing model. Investigations examined cost-performance trade-offs on serverless and cluster-centric ETL solutions. the fact that serverless approaches result in reduced costs for idle resources, but tend to be more expensive than non-serverless architectures under sustained workloads, emphasizing the need of workload-aware cost models [14].

10.48047/jocaaa.2024.33.08.394

In Scalability Eclipse: Truths and Myths in Cloud Native ETL tools, on the one hand we see that linear scaling is not guaranteed because of network bottlenecks, data skew and orchestration overhead. Indeed, previous studies indicated that efficient parallelism and good workload division are pivotal factors in achieving scalability efficiency when processing terabyte-scale data on distributed environments [15].

Trade-offs between ease-of-use and performance-control in comparison of managed ETL services with open-source frameworks. Managed service was better than open source in terms of simplified deployment and fewer operational challenges, but at the cost of being less flexible towards optimization. These studies stressed that tools should be selected based on organizational skills and performance needs, too [16].

More recent research included real-time and near-real-time ETL situations, which is indicative of increased interest in low-concurrency analytics. These studies investigated stream-based ETL pipelines and the hybrid batch-stream processing approach. The study also suggested that traditional batch-based ETL tools would be insufficient to meet low-latency requirements and event-driven architectures would need to be adopted [17].

Framework proposals from the literature were created for reproducibility of evaluation of ETL tools. These were standardized datasets, metrics and operating conditions to enable fair comparison. However, they noted inadequacy in adapting to changes of cloud services and different pricing schemas among providers [18].

Security and reliability concerns were also discussed in relation to cloud-native ETL pipelines. Papers had joint topics such: Data privacy, Access control and Fault tolerance. In cloud platforms, security mechanisms are available but improper configuration and the dependence on managed services were recognized as potential threats to the robustness of pipelines [19].

Recent studies have also focused on integration of ETL pipelines and analytics and machine learning workflows. These works demonstrated the necessity of building scalable and cost-effective data preparation layers that can efficiently support iterative processing and compute-intensive workloads. As a whole, the literature speaks to a need for more comprehensive benchmarking efforts that involve performance, cost and scalability with all other parameters held constant across experiments [20].

3. Methodology

This article employs a quantitative experimental approach to compare the performance, cost efficiency and scalability of cloud-native ETL systems in a controlled environment. All experiments were conducted in a leveled cloud-based environment, so that infrastructure bias does not take place. Equal virtual machine configuration, storage type and network bandwidth were provisioned to each of the ETL tool for being fair. We used different-size datasets—including small (10 GB) and large (2TB)—to demonstrate behavior under various workload strengths.




Figure 1. Flowchart of the proposed benchmarking methodology for cloud-native ETL tools, illustrating performance, cost, and scalability evaluation steps.

Figure 1. depicts the detailed flowchart of the proposed benchmarking process for assessing cloud-native ETL tools. The first step consists in selecting some benchmark datasets with different sizes and types, then creating a homogeneous cloud infrastructure to make possible an equal comparison. For each ETL tool, it is finally deployed with the predefined configurations and executed on the data in order to time stamp its execution. Performance values (e.g., execution time and throughput) are derived from the stored values. Cloud billing data is then analyzed for the calculation of total cost and a proper to-scale value with respect to each amount unit. We create a scaling chart by further increasing the level of parallelism and recording run time at scaled settings. Finally, the approach calculates scalability and global efficiency ratios that allow to perform comparative evaluation of ETL engines on performance, cost effectiveness and parallelizability.

Performance Analysis was mainly performed in execution efficiency and data processing capacity. The total running time of an ETL job was calculated by the difference between end time and start time (in second) represented as:

$$T_{exec} = T_{end} - T_{start} \quad (1)$$

This metric captures the end-to-end latency of data ingestion, transformation, and loading. To further quantify processing efficiency, throughput was calculated as the ratio of processed data volume to execution time, enabling comparative assessment across tools and dataset sizes, given by

$$Throughput = \frac{D_{size}}{T_{exec}} \quad (2)$$

Cost analysis was performed using actual cloud billing data collected during pipeline execution. The total cost incurred by an ETL job was modeled as the product of resource usage duration and the unit cost of cloud resources. This relationship is expressed as

$$C_{total} = \sum_{i=1}^n R_i \times U_i \quad (3)$$

where R_i represents the runtime of the i^{th} resource and U_i denotes its corresponding unit price. To enable fair cost comparison across varying dataset sizes, cost efficiency was normalized as cost per terabyte processed, defined as

$$C_{norm} = \frac{C_{total}}{D_{size}} \quad (4)$$

Scalability was evaluated by progressively increasing the level of parallelism and observing corresponding changes in performance. The scalability factor was computed as the ratio of execution time under single-instance execution to execution time under scaled conditions, as shown in

$$S = \frac{T_1}{T_k} \quad (5)$$

where T_1 is the time needed to execute using one processing unit and T_k is common runtime with k parallel units. This measure indicates how well an ETL tool uses extraneous capacity.

Lastly, global efficiency was measured by combining performance and cost factors into a single efficiency score. This can be seen as a comprehensive performance index of the tool and is expressed as,

$$E = \frac{Throughput}{C_{norm}} \quad (6)$$

Higher values of this index indicate superior performance-cost balance. Together, these methodological steps and analytical models provide a comprehensive and reproducible framework for benchmarking cloud-native ETL tools across performance, cost, and scalability dimensions.

4. Results and Discussion

This section presents and analyzes the experimental results obtained from benchmarking cloud-native ETL tools in terms of performance, cost efficiency, and scalability. The evaluation was

conducted using identical cloud infrastructure and standardized workloads to ensure fairness and reproducibility.

Performance Analysis

The capacities that have been observed for the execution time and throughput are summarized in Table 1. As reported in the table, the serverless ETL tool (Tool A) had the minimum execution time (42 minutes) and maximum throughput (24.1 GB/min). This higher performance is due to its elastic resource allocation and less infrastructure management overhead. In comparison, the managed cluster-based ETL tool (Tool B) completed in 55 min and open-source ETL framework (Tool C) achieved longest execution time of 68 min which representing higher orchestration and resource management overhead.




Figure 2. Execution Time Comparison Across Cloud-Native ETL Tools.

Figure 2 visually compares execution times across ETL tools, clearly demonstrating the performance superiority of serverless architectures for medium-scale workloads. The results indicate that execution efficiency varies significantly depending on the underlying execution model and resource management strategy.

Table 1. Performance Comparison of Cloud-Native ETL Tools

ETL Tool	Execution Time (min)	Throughput (GB/min)
Tool A (Serverless)	42	24.1
Tool B (Managed Cluster)	55	18.4
Tool C (Open Source)	68	14.7

Cost Efficiency Evaluation

In Table 2 we present our cost analysis results, which is the cost normalized by terabyte processed. Tool A had the lowest cost at 18.6 USD/TB, which indicated that serverless ETL was a good choice for inconsistent and varying workloads. Tool B was the most expensive due to keeping resources reserved, and Tool C proved a moderate cost-effective solution (but added operational overhead).




Figure 3. Cost per Terabyte Comparison Across Cloud-Native ETL Tools

The cost per TB comparison (Figure 3) shows that those which support the dynamic allocation of resources have a better tradeoff between cost and performance. It is concluded to give consideration for workload-aware ETL tool selection, especially under the pay-as-you-go cloud pricing schemes.

Scalability and Overall Efficiency

Scalability Results The scalability results show that Tool B has the highest S, equal to 0.97 which suggests a nearly linear scaling with increased parallelism. Tool A was close with a scalability factor of 0.94 and Tool C showed slower scalability as a result of synchronization and splitting data overheads.




Figure 4. Scalability Factor Comparison of Cloud-Native ETL Tools

Figure 4 presents the scalability comparison, demonstrating that managed cluster-based solutions are more stable under sustained large-scale workloads. The overall efficiency index, combining performance and cost metrics, favored Tool A with the highest efficiency score of 1.30, indicating superior overall effectiveness.

Table 2. Cost, Scalability, and Efficiency Metrics

ETL Tool	Cost per TB (USD)	Scalability Factor (S)	Efficiency Index (E)
Tool A (Serverless)	18.6	0.94	1.30

Tool B (Managed Cluster)	24.9	0.97	0.74
Tool C (Open Source)	21.3	0.88	0.69

Discussion Summary

In summary, the results show that there is no single ETL tool that performs best in all aspects. Serverless ETL tools work well for performance- and cost-sensitive workloads, but managed cluster-based solutions provide better scalability of throughput if you have large amounts of data to process continuously. Open-source ETL frameworks are flexible but need to be carefully configured in order to obtain good performance. We offer these insights to help others focus on the selection of cloud-native ETL tools, with regard to both workload and organizational priority.

5. Conclusion

We provided a thorough benchmarking in the context of controlled experiments of cloud-native ETL tools in terms of performance, cost-effectiveness, and scalability. The study enabled the comparison of serverless, managed cluster-based and open-source ETL solutions on a common cloud infrastructure under fair conditions, providing impartial and repeatable results on their operational characteristics. We showed that serverless ETL tools provide higher execution efficiency and lower cost per terabyte for small to medium-sized workloads because of elastic resource allocation and little infrastructure overhead. The ETL tools implemented on managed cluster platform show good scale-out features and achieve near-linear throughput performance when running at higher degree of parallelism, therefore capable for continuous big data processing. In comparison, opensource ETL frameworks provided flexibility and was moderately cost-effective with careful tuning to avoid performance degradation and scaling limitations.

The overall efficiency index showed that the performance and cost should be considered together rather than independently for choosing an ETL solution. The results indicate that workload properties, amount of data processed and rate at which is done so are all key drivers paving the way in choosing the most suitable cloud-native ETL architecture. In summary, we believe that this benchmarking process offers useful guidance to data engineers and architects and informs them about the tool choices. In future, we will also consider the real-time streaming workloads, multi-

cloud environment and energy efficiency metrics to make cloud-native ETL evaluation more robust.

References

1. Guo, Y.; Varbanescu, A.L.; Iosup, A.; Martella, C.; Willke, T.L. Benchmarking graph-processing platforms: A vision. In Proceedings of the 5th ACM/SPEC International Conference on Performance Engineering, Dublin, Ireland, 22–26 March 2014; pp. 289–292. [Google Scholar]
2. Dayarathna, M.; Suzumura, T. Benchmarking graph data management and processing systems: A survey. *arXiv* **2020**, arXiv:2005.12873. [Google Scholar]
3. Uta, A.; Au, S.; Ilyushkin, A.; Iosup, A. Elasticity in graph analytics? A benchmarking framework for elastic graph processing. In Proceedings of the 2018 IEEE International Conference on Cluster Computing (CLUSTER), Belfast, UK, 10–13 September 2018; pp. 381–391. [Google Scholar]
4. Jiang, J.; Xiao, P.; Yu, L.; Li, X.; Cheng, J.; Miao, X.; Zhang, Z.; Cui, B. PSGraph: How Tencent trains extremely large-scale graphs with Spark? In Proceedings of the 2020 IEEE 36th International Conference on Data Engineering (ICDE), Dallas, TX, USA, 20–24 April 2020; pp. 1549–1557. [Google Scholar]
5. Guo, R.; Zhao, Y.; Zou, Q.; Fang, X.; Peng, S. Bioinformatics applications on Apache Spark. *GigaScience* **2018**, 7, giy098. [Google Scholar] [CrossRef] [PubMed]
6. Mohammed, W.M.S.; Maa, A.K.J. An efficient approach to extract and store big semantic web data using Hadoop and Apache Spark GraphX. *ADCAIJ Adv. Distrib. Comput. Artif. Intell. J.* **2024**, 13, e31506. [Google Scholar]
7. Zhou, H.; Sun, G.; Fu, S.; Wang, L.; Hu, J.; Gao, Y. Internet financial fraud detection based on a distributed big data approach with Node2vec. *IEEE Access* **2021**, 9, 43378–43386. [Google Scholar] [CrossRef]
8. Nazari, E.; Shahriari, M.H.; Tabesh, H. BigData analysis in healthcare: Apache hadoop, apache spark and apache flink. *Front. Health Inform.* **2019**, 8, 14. [Google Scholar] [CrossRef]

10.48047/jocaaa.2024.33.08.394

9. He, C.; Huang, Y.; Wang, C.; Wang, N. Dynamic data partitioning strategy based on heterogeneous Flink cluster. In Proceedings of the 2022 5th International Conference on Artificial Intelligence and Big Data (ICAIBD), Chengdu, China, 27–30 May 2022; pp. 355–360. [Google Scholar] [CrossRef]
10. Rabl, T.; Traub, J.; Katsifodimos, A.; Markl, V. Apache Flink in current research. *it-Inf. Technol.* **2016**, *58*, 157–165. [Google Scholar] [CrossRef]
11. Best Practices—Airflow Documentation. Available online: <https://airflow.apache.org/docs/apache-airflow/stable/best-practices.html> (accessed on 8 August 2022).
12. Aquilanti, P.-Y.; Kendrex, S.; Koop, M. AWS Batch Dos and Don'ts: Best Practices in a Nutshell|AWS HPC Blog. Available online: <https://aws.amazon.com/blogs/hpc/aws-batch-best-practices/> (accessed on 8 August 2022).
13. Liston, B. Creating a Simple 'Fetch & Run' AWS Batch Job|AWS Compute Blog. Available online: <https://aws.amazon.com/blogs/compute/creating-a-simple-fetch-and-run-aws-batch-job/> (accessed on 8 August 2022).
14. VPS Serveriai—Interneto Vizija. Available online: <https://www.iv.lt/vps-serveriai/#kontaineriai> (accessed on 10 August 2022).
15. Amazon EC2 Pricing—Amazon Web Services. Available online: <https://aws.amazon.com/ec2/pricing/> (accessed on 10 August 2022).
16. Khalid, M.; Yousaf, M.M. A Comparative Analysis of Big Data Frameworks: An Adoption Perspective. *Appl. Sci.* **2021**, *11*, 11033. [Google Scholar] [CrossRef]
17. Rovnyagin, M.M.; Shipugin, V.A.; Ovchinnikov, K.A.; Durachenko, S.V. Intelligent container orchestration techniques for batch and micro-batch processing and data transfer. *Procedia Comput. Sci.* **2021**, *190*, 684–689. [Google Scholar] [CrossRef]
18. Pishgoo, B.; Azirani, A.A.; Raahemi, B. A hybrid distributed batch-stream processing approach for anomaly detection. *Inf. Sci.* **2020**, *543*, 309–327. [Google Scholar] [CrossRef]

10.48047/jocaaa.2024.33.08.394

19. ISO/IEC TR 30172:2023; Internet of THINGS (IoT)—Digital twin—Use Cases. International Organization for Standardization: Geneva, Switzerland, 2023. Available online: <https://www.iso.org/standard/81578.html> (accessed on 24 January 2024). ISO/IEC 30173:2023; Digital Twin—Concepts and Terminology. International Organization for Standardization: Geneva, Switzerland, 2023. Available online: <https://www.iso.org/standard/81442.html> (accessed on 24 January 2024). Wang, Z.;
20. Gupta, R.; Han, K.; Wang, H.; Ganlath, A.; Ammar, N.; Tiwari, P. Mobility Digital Twin: Concept, Architecture, Case Study, and Future Challenges. IEEE Internet Things J. 2022, 9, 17452–17467. [Google Scholar] [CrossRef]