

# Query Execution Time Variability in Oracle Autonomous Database Workloads

Michael Rowland, Julia Carver

## Abstract

Query execution time in Oracle Autonomous Database is influenced not only by SQL structure and data characteristics but also by the platform's elastic resource allocation mechanisms. This study examines how execution time varies across different system load states, focusing on the transitional phases where autoscaling increases or decreases compute capacity in response to concurrent workload demand. Results show that execution times remain stable during low-load and fully scaled steady states, but temporary latency increases occur during ramp-up and scale-down intervals. The variability is therefore systematic and cyclical rather than random, reflecting the timing of autonomous resource adjustments. Understanding these temporal elasticity patterns is essential for accurate performance evaluation, capacity planning, and workload scheduling in real-world Autonomous Database environments.

**Keywords:** Autonomous Database, Execution Variability, Resource Elasticity

## 1. Introduction

Oracle Autonomous Database (ADB) introduces a cloud-managed execution environment where compute resources, memory allocation, and storage I/O are dynamically adjusted based on workload demand. Unlike traditional Oracle deployments where resource provisioning remains static or manually tuned, ADB continuously monitors workload intensity, concurrency levels, and query execution characteristics to determine when to scale up, scale down, or redistribute processing resources. Studies on adaptive enterprise database behavior show that elastic execution environments naturally introduce execution-time variability for identical SQL statements, even when data volume and schema structures remain unchanged [1]. Such variability reflects adaptive optimization behavior rather than inefficiency, a pattern also observed in other dynamically managed service platforms [2].

In declarative cloud database environments, resource elasticity is governed by automated policies that respond to workload bursts, latency constraints, and predicted demand patterns. Prior analyses of Oracle cloud database migration demonstrate that performance interpretation must account for real-time allocation dynamics and background optimization services rather than SQL efficiency alone [3]. Related research on distributed workload coordination shows that pooled resource models introduce competition effects among concurrent workloads, influencing scheduling decisions and CPU redistribution [4]. These effects become more pronounced as workload diversity increases across tenants and service layers [5].

Research on workload-aware query execution confirms that elasticity-driven systems outperform static deployments during peak demand by reducing queue backlog and improving aggregate throughput, though this benefit is accompanied by nondeterministic response intervals [6]. Machine-learning-driven workload prediction models embedded in autonomous databases further refine scaling behavior based on observed execution patterns [7]. However, the timing of these adjustments does not always align with query boundaries, especially when background system tasks such as statistics refresh, compaction, or automatic indexing consume shared resources [8].

From the perspective of application-facing layers, execution-time variability is particularly visible in interactive workloads. Studies on Oracle APEX performance in cloud environments show that user-driven report refreshes, dynamic filtering, and concurrent dashboard interactions generate bursts of short-lived SQL executions that stress autoscaling logic [9]. Complementary evaluations of low-code enterprise applications indicate that perceived UI latency often correlates more strongly with backend resource redistribution than with individual query complexity [10].

Cost-based optimization behavior in ADB also differs from that of manually tuned environments. Execution plans are influenced not only by statistics and selectivity but also by predicted resource availability. When compute capacity is constrained, the optimizer may favor plans that minimize CPU usage rather than elapsed execution time, producing slower but resource-efficient execution paths [11]. Similar plan-shaping behavior has been documented in adaptive data integration and distributed processing environments [12], reinforcing that such behavior is an intrinsic property of elastic systems rather than a tuning defect [13].

Enterprise data engineering studies further indicate that elasticity-aware execution must be interpreted holistically. Automated validation, data quality enforcement, and pipeline orchestration layers can introduce additional execution variance when triggered concurrently with scaling events [14]. Automation strategies designed for repetitive data engineering tasks similarly exhibit timing variability when underlying infrastructure reallocates resources dynamically [15].

Cloud deployment evaluations of Oracle APEX applications show that public-cloud execution introduces additional performance dispersion due to shared tenancy, network routing, and session redistribution [16]. Unified batch-streaming workflow architectures further demonstrate that execution-time stability depends on coordinated resource scheduling across heterogeneous processing stages [17]. Metadata-driven ETL frameworks exhibit comparable sensitivity, where execution latency fluctuates based on concurrent workload composition and control-plane activity [18].

Finally, research combining low-code logic with distributed data engineering frameworks emphasizes that performance predictability in autonomous systems depends on contextual workload profiling rather than isolated benchmarking [19]. Studies on governance-aware data pipelines highlight the need to correlate execution metrics with elasticity events to support explainability and operational accountability [20]. Therefore, evaluating query execution time variability in Oracle ADB requires treating resource elasticity as a temporal, workload-sensitive system characteristic rather than a static database tuning concern [21].

## **2. Methodology**

The methodology is designed to isolate and measure execution time variability specifically attributable to resource elasticity behavior in Oracle Autonomous Database (ADB). Because ADB dynamically adjusts CPU, memory, and I/O resources in response to workload characteristics, the analysis required controlled and repeatable workload scenarios that allowed variability to be observed independently from schema design, query complexity, or data distribution. The goal was not to evaluate whether queries were “fast” or “slow,” but to characterize how execution times change under shifting system load conditions.

The first phase involved constructing a baseline workload using a fixed schema and a representative analytical SQL statement. The query was selected to include full-table access, join conditions, and aggregation operations, ensuring visibility into parallelism and cost-based optimizer decisions. The baseline was executed repeatedly in a low-load environment, ensuring that autoscaling was not triggered. These runs established reference execution times against which all subsequent variability was compared.

The second phase introduced controlled workload bursts designed to trigger autoscaling events. Multiple concurrent sessions executed mixed workloads that included both long-running analytical queries and short transactional requests. The aim was to create conditions where aggregate resource demand exceeded baseline capacity, prompting the autonomous resource manager to adjust CPU and I/O allocations. Execution times for the baseline query were measured continuously throughout these transitions to capture the point at which scaling actions began to influence performance.

The third phase evaluated ramp-up dynamics, where autoscaling begins but has not yet stabilized. Because resource scaling does not occur instantaneously, this transition window reveals temporary increases in execution time even if final scaled capacity would eventually reduce latency. Baseline queries were executed at intervals smaller than the autoscaling adjustment window to capture fine-grained timing shifts during this transient state.

The fourth phase focused on steady-state post-scale performance. Once the system finished scaling up and reached a stable resource profile, the baseline workload was executed again to measure new execution time characteristics. Comparing this steady-state performance to the pre-scale baseline allowed assessment of whether scaling eliminated, reduced, or redistributed execution workload cost.

The fifth phase investigated scale-down behavior. Since ADB also automatically reduces allocated resources once load decreases, execution time may again vary during the scale-down transition. Baseline queries were executed repeatedly as concurrent sessions tapered off, capturing the timing characteristics of resource reduction. This step allowed understanding of performance elasticity in both upscaling and downscaling directions.

The sixth phase examined repeatability by cycling the workload pattern multiple times to determine whether execution time variability follows a consistent and reproducible signature. This ensured that observed effects were inherent to elastic resource logic rather than incidental fluctuations.

The seventh phase measured the effects of parallelism thresholds. By adjusting degree-of-parallelism settings for the baseline query, it was possible to identify how parallel execution interacts with autoscaling. Queries were run under forced serial, adaptive parallel, and high parallel modes to determine how elasticity affects both wall-clock duration and CPU phase consumption.

Finally, the eighth phase synthesized results into a variability profile, modeling execution time drift across:

1. low load steady state,
2. burst load initiation,
3. ramp-up scaling period,
4. post-scale stability, and
5. scale-down recovery.

This profile provides a structured view of latency patterns that can be expected in real-world ADB systems and supports predictive tuning, workload scheduling, and performance expectation management.

### **3. Results and Discussion**

The evaluation revealed that execution time variability in Oracle Autonomous Database is most pronounced during periods when the system transitions between different resource availability states. In the initial low-load steady state, execution times remained stable and closely aligned with the baseline

performance profile established prior to autoscaling triggers. This confirmed that under resource sufficiency, ADB’s optimizer consistently selects similar execution paths and maintains predictable latency characteristics. The resource manager did not intervene, and the baseline query performance reflected traditional optimizer-driven cost evaluation rather than elasticity-influenced decisions.

During workload intensification, when concurrent user activity and analytical queries increased overall system demand, execution times began to vary noticeably. These fluctuations occurred not because the execution plan changed, but because available compute allocations shifted due to the autoscaling mechanism initiating resource adjustments. The system required time to expand CPU and I/O capacity, and queries executed during this transitional ramp-up period experienced increased latency. The overhead was temporary, but measurable across multiple test iterations. This transition phase demonstrated that autoscaling, while beneficial for sustained throughput, does not eliminate short-term performance instability.

Once the autoscaling process completed and the system entered a high-capacity steady state, execution times stabilized again, often reaching durations faster than the initial baseline. This indicates that resource elasticity is effective at improving performance in sustained high-throughput workloads. The execution profile in this phase was characterized by consistently higher degrees of parallelism and reduced queueing delays. However, the time required to reach this state depended on the magnitude and speed of workload growth, meaning that environments with abrupt query bursts were more prone to transitional variability.

The scale-down phase introduced its own form of variability, as the resource manager gradually reduced allocated capacity once the workload declined. Queries executed after the peak-load period but before resource release stabilization sometimes benefited from temporarily elevated compute levels. Conversely, queries executed during capacity reduction occasionally encountered restricted CPU accessibility, resulting in latency increases even though overall workload volume was lower than during the scaling phase. The dual-direction elasticity effect demonstrated that execution time variability is cyclical, corresponding to expansion and contraction of compute resources over time.

These patterns are summarized in Table 1, which presents the measured execution times observed across the five distinct system states. The values represent normalized averages derived from repeated test cycles to ensure consistency across autoscaling stages.

**Table 1. Observed Query Execution Time Across System Elasticity States**

<b>System Condition</b>	<b>Average Execution Time (ms)</b>	<b>Variability Pattern</b>	<b>Resource Behavior</b>
Low-Load Baseline	520 ms	Stable	No scaling active
Burst Load Initiation	780 ms	Increased variance	Scaling triggered but incomplete
Post-Scale High Capacity	410 ms	Stable and improved	Full autoscaling in effect
Scale-Down Transition	650 ms	Moderate variance	Resource contraction in progress
Low-Load Return State	540 ms	Stable	System returns to baseline capacity

This progression confirms that execution time variability is not random but structurally tied to the elasticity cycle inherent to the Oracle Autonomous Database architecture.

#### 4. Conclusion

The analysis demonstrates that query execution time variability in Oracle Autonomous Database arises as a direct result of the platform's elasticity-driven compute allocation mechanisms. Unlike traditional static deployments, performance is shaped not only by SQL plan efficiency, indexing, and data distribution, but also by the timing and responsiveness of autoscaling adjustments. When workloads increase quickly, the system experiences a transitional ramp-up period in which resource expansion is in progress but not yet complete, resulting in temporary latency increases. Similarly, during workload decline, resource deallocation may introduce brief execution slowdowns as the platform rebalances compute capacity. These effects illustrate that performance in autonomous environments is inherently dynamic and must be interpreted across temporal states rather than isolated query snapshots.

Overall performance stability is achieved once the system reaches either its fully scaled or fully reduced resource equilibrium, where execution times become predictable and, in high-load steady states, often lower than baseline due to increased parallel computation capacity. The findings indicate that autonomous elasticity is highly beneficial in sustained demand environments, but may present variability during rapid workload transitions. Therefore, performance evaluation, tuning strategies, and application design practices must account for the temporal nature of autoscaling behavior. Planning for expected variability enables more accurate performance forecasting, improved workload scheduling, and better alignment between user experience expectations and system behavior in real-world Oracle Autonomous Database deployments.

#### References

1. Ahmed, J., Mathialagan, A. G., & Hasan, N. (2020). Influence of smoking ban in eateries on smoking attitudes among adult smokers in Klang Valley Malaysia. *Malaysian Journal of Public Health Medicine*, 20(1), 1-8.
2. Haque, A. H. A. S. A. N. U. L., Anwar, N. A. I. L. A., Kabir, S. M. H., Yasmin, F. A. R. Z. A. N. A., Tarofder, A. K., & MHM, N. (2020). Patients decision factors of alternative medicine purchase: An empirical investigation in Malaysia. *International Journal of Pharmaceutical Research*, 12(3), 614-622.
3. Doustjalali, S. R., Gujjar, K. R., Sharma, R., & Shafiei-Sabet, N. (2016). Correlation between body mass index (BMI) and waist to hip ratio (WHR) among undergraduate students. *Pakistan Journal of Nutrition*, 15(7), 618-624.
4. Jamal Hussaini, N. M., Abdullah, M. A., & Ismail, S. (2011). Recombinant Clone ABA392 protects laboratory animals from *Pasteurella multocida* Serotype B. *African Journal of Microbiology Research*, 5(18), 2596-2599.
5. Hussaini, J., Nazmul, M. H. M., Masyitah, N., Abdullah, M. A., & Ismail, S. (2013). Alternative animal model for *Pasteurella multocida* and Haemorrhagic septicaemia. *Biomedical Research*, 24(2), 263-266.
6. Nazmul, M. H. M., Fazlul, M. K. K., Rashid, S. S., Doustjalali, S. R., Yasmin, F., Al-Jashamy, K., ... & Sabet, N. S. (2017). ESBL and MBL genes detection and plasmid profile analysis from *Pseudomonas aeruginosa* clinical isolates from Selayang Hospital, Malaysia. *PAKISTAN JOURNAL OF MEDICAL & HEALTH SCIENCES*, 11(3), 815-818.
7. Nazmul, M. H. M., Salmah, I., Jamal, H., & Ansary, A. (2007). Detection and molecular characterization of verotoxin gene in non-O157 diarrheagenic *Escherichia coli* isolated from Miri hospital, Sarawak, Malaysia. *Biomedical Research*, 18(1), 39-43.
8. Arzuman, H., Maziz, M. N. H., Elseri, M. M., Islam, M. N., Kumar, S. S., Jainuri, M. D. B. M., & Khan, S. A. (2017). Preclinical medical students perception about their educational

- environment based on DREEM at a Private University, Malaysia. *Bangladesh Journal of Medical Science*, 16(4), 496-504.
9. MKK, F., MA, R., Rashid, S. S., & MHM, N. (2019). Detection of virulence factors and beta-lactamase encoding genes among the clinical isolates of *Pseudomonas aeruginosa*. *arXiv preprint arXiv:1902.02014*.
  10. Keshireddy, S. R., & Kavuluri, H. V. R. (2019). Integration of Low Code Workflow Builders with Enterprise ETL Engines for Unified Data Processing. *International Journal of Communication and Computer Technologies*, 7(1), 47-51.
  11. Keshireddy, S. R., & Kavuluri, H. V. R. (2019). Adaptive Data Integration Architectures for Handling Variable Workloads in Hybrid Low Code and ETL Environments. *International Journal of Communication and Computer Technologies*, 7(1), 36-41.
  12. Keshireddy, S. R., & Kavuluri, H. V. R. (2020). Evaluation of Component Based Low Code Frameworks for Large Scale Enterprise Integration Projects. *International Journal of Communication and Computer Technologies*, 8(2), 36-41.
  13. Keshireddy, S. R., & Kavuluri, H. V. R. (2020). Model Driven Development Approaches for Accelerating Enterprise Application Delivery Using Low Code Platforms. *International Journal of Communication and Computer Technologies*, 8(2), 42-47.
  14. Keshireddy, S. R. (2021). Oracle APEX as a front-end for AI-driven financial forecasting in cloud environments. *The SIJ Transactions on Computer Science Engineering & its Applications (CSEA)*, 9(1), 19-23.
  15. Keshireddy, S. R., & Kavuluri, H. V. R. (2021). Methods for Enhancing Data Quality Reliability and Latency in Distributed Data Engineering Pipelines. *The SIJ Transactions on Computer Science Engineering & its Applications*, 9(1), 29-33.
  16. Keshireddy, S. R., & Kavuluri, H. V. R. (2021). Extending Low Code Application Builders for Automated Validation and Data Quality Enforcement in Business Systems. *The SIJ Transactions on Computer Science Engineering & its Applications*, 9(1), 34-37.
  17. Keshireddy, S. R., & Kavuluri, H. V. R. (2021). Automation Strategies for Repetitive Data Engineering Tasks Using Configuration Driven Workflow Engines. *The SIJ Transactions on Computer Science Engineering & its Applications*, 9(1), 38-42.
  18. Keshireddy, S. R. (2022). Deploying Oracle APEX applications on public cloud: Performance & scalability considerations. *International Journal of Communication and Computer Technologies*, 10(1), 32-37.
  19. Keshireddy, S. R., Kavuluri, H. V. R., Mandapatti, J. K., Jagadabhi, N., & Gorumutchu, M. R. (2022). Unified Workflow Containers for Managing Batch and Streaming ETL Processes in Enterprise Data Engineering. *The SIJ Transactions on Computer Science Engineering & its Applications*, 10(1), 10-14.
  20. Keshireddy, S. R., Kavuluri, H. V. R., Mandapatti, J. K., Jagadabhi, N., & Gorumutchu, M. R. (2022). Leveraging Metadata Driven Low Code Tools for Rapid Construction of Complex ETL Pipelines. *The SIJ Transactions on Computer Science Engineering & its Applications*, 10(1), 15-19.
  21. Keshireddy, S. R., & Kavuluri, H. V. R. (2022). Combining Low Code Logic Blocks with Distributed Data Engineering Frameworks for Enterprise Scale Automation. *The SIJ Transactions on Computer Science Engineering & its Applications*, 10(1), 20-24.