



# LEVERAGING AI TO IMPROVE PERFORMANCE TUNING IN POST-MIGRATION ORACLE CLOUD ENVIRONMENTS

**Prasad Manda**

Principal Database Engineer/Architect, 3M Company/Solventum, USA.

## ABSTRACT

*As enterprises increasingly migrate Oracle workloads to cloud platforms like Oracle Cloud Infrastructure (OCI) and Amazon Web Services (AWS), performance degradation post-migration has emerged as a critical concern. Despite the promised scalability and agility of cloud environments, variations in infrastructure, suboptimal configurations, and outdated performance tuning strategies can result in inefficient workloads and increased operational costs. Traditional methods such as manual AWR analysis, SQL plan management, and reactive diagnostics often fail to scale or adapt to the dynamic nature of cloud-native architectures.*

*This paper explores how Artificial Intelligence (AI) is redefining performance tuning in post-migration Oracle environments. It highlights the role of AI-powered features—including autonomous indexing, predictive scaling, anomaly detection, and automated root cause analysis—in addressing tuning challenges across cloud-hosted Oracle databases. We analyze the integration of AI capabilities within Oracle's ecosystem, such as the Autonomous Database and Oracle Management Cloud, as well as third-party AIOps tools. Through comparative analysis, architectural models, and practical guidance, we demonstrate how AI not only enhances the DBA's ability to*

*maintain high-performance systems but also shifts their role from reactive firefighting to proactive performance engineering.*

*The findings underscore that AI-driven performance tuning is not a replacement for database administrators, but rather an enabler of strategic optimization in complex, distributed cloud landscapes.*

**Keywords:** Oracle Cloud Infrastructure (OCI) , Performance Tuning, Cloud Migration, Artificial Intelligence (AI), Autonomous Database, AWR and ASH Analytics, SQL Plan Optimization, AI-Driven Database Management

**Cite this Article:** Prasad Manda. (2023). Leveraging AI to Improve Performance Tuning in Post-Migration Oracle Cloud Environments. *International Journal of Artificial Intelligence & Machine Learning (IJAIML)*, 2(1), 235-247.

DOI: [https://doi.org/10.34218/IJAIML\\_02\\_01\\_021](https://doi.org/10.34218/IJAIML_02_01_021)

---

## 1. State of the Art in Oracle Performance Management

### 1.1 Cloud Migration and Oracle Performance Trends

Enterprise adoption of cloud infrastructure has accelerated significantly, with Oracle workloads forming a key part of digital transformation strategies. According to Gartner and IDC forecasts, more than 70% of enterprise Oracle databases are expected to transition to cloud platforms—such as Oracle Cloud Infrastructure (OCI), Oracle Exadata Cloud@Customer, and AWS EC2 with Oracle Enterprise Edition—by 2025. While cloud migration offers clear advantages in scalability, elasticity, and cost optimization, it also introduces new tuning challenges due to abstracted infrastructure, multi-tenant execution environments, and on-demand compute models.

### 1.2 Limitations of Traditional Performance Tuning Techniques

Historically, Oracle Database Administrators (DBAs) have relied on a suite of mature, yet manually intensive, tools and methodologies for performance optimization:

- **AWR (Automatic Workload Repository)** and **ASH (Active Session History)** for time-based workload analytics,
- **SQL Plan Management (SPM)** and optimizer hints to guide execution plans,
- Manual **index tuning, partitioning, and query rewrites** for performance bottlenecks.

These approaches are effective in predictable, static environments. However, they are increasingly insufficient in dynamic, cloud-native ecosystems where workloads scale

elastically, telemetry changes frequently, and tuning windows are narrow or nonexistent. Moreover, traditional methods remain largely reactive and depend heavily on DBA experience and heuristics.

### ***1.3 AI-Driven Innovations in Oracle Performance Tuning***

The integration of Artificial Intelligence (AI) and Machine Learning (ML) into database performance management is reshaping the Oracle tuning landscape. Notable developments include:

- **Oracle Autonomous Database**, which uses AI for auto-indexing, adaptive query plans, patch automation, and self-tuning without DBA intervention.
- **Oracle Management Cloud (OMC)**, offering AI/ML-based anomaly detection, cross-domain correlation, and proactive alerting based on real-time telemetry.
- **Real Application Testing (RAT) and SQL Performance Analyzer (SPA)**, which become more powerful when augmented with predictive analytics and pattern recognition capabilities.
- **Third-party AIOps tools** (e.g., Dynatrace, Datadog, AppDynamics) that integrate with Oracle and apply AI to detect, diagnose, and resolve performance anomalies in real time.

These tools shift the paradigm from “observe and act” to “predict and automate,” enhancing tuning accuracy, reducing mean time to resolution (MTTR), and freeing DBAs from low-level diagnostics.

### ***1.4 Research and Industry Gaps***

Despite these innovations, several gaps remain:

- Most AI-based tuning systems are evaluated in isolation, with limited benchmarking in multi-cloud or hybrid Oracle environments.
- There is a lack of open-source research or benchmarking datasets to evaluate AI tuning effectiveness in real-world, post-migration scenarios.
- End-to-end correlation across the stack (network, storage, database, application) is still fragmented in many toolchains, hindering holistic root cause analysis.
- Adoption barriers exist due to organizational inertia, trust issues in autonomous decisions, and the need for fine-grained control over mission-critical SQL performance.

## 2. Post-Migration Performance Challenges in Oracle Cloud Environments

The transition of Oracle workloads from on-premises infrastructure to cloud platforms such as Oracle Cloud Infrastructure (OCI) and Amazon Web Services (AWS) introduces a series of complex performance management issues. These challenges often arise despite successful functional migration, and they can undermine cloud benefits if not properly addressed. This section identifies and analyzes the key technical hurdles encountered in the post-migration phase.

### 2.1 Hardware and Platform Disparities

One of the most immediate causes of performance degradation post-migration is the difference in underlying hardware and infrastructure characteristics. On-premises Oracle deployments often run on optimized platforms such as Oracle Exadata, which provide advanced features like Smart Scan and Hybrid Columnar Compression. When these workloads are migrated to commodity cloud instances (e.g., AWS EC2 or standard OCI VM shapes), performance characteristics change due to:

- Lower IOPS throughput,
- Lack of hardware-level query offloading,
- Variability in compute and memory configurations.

This mismatch can result in slower query execution times, increased wait events, and degraded throughput.

### 2.2 Storage and I/O Bottlenecks

Cloud environments often involve shared or virtualized storage systems where performance depends on factors such as:

- Provisioned IOPS or storage tiers,
- File system configurations (e.g., XFS vs. ASM),
- Network latency in accessing block/object storage.

Insufficient I/O tuning or inappropriate storage class selection may lead to bottlenecks in read/write operations, particularly for data-intensive OLTP or hybrid workloads.

### 2.3 Suboptimal Cloud-Native Configurations

Many post-migration environments suffer from poorly tuned cloud-native parameters. These include:

- Inappropriate VM shapes or autoscaling policies,
- Incorrect database block sizes,
- Misaligned PGA/SGA settings in relation to new compute limits.

Default configurations often do not reflect workload-specific characteristics, leading to inefficient memory usage, excessive context switching, and increased disk sorts or temporary tablespace usage.

### 2.4 Workload Behavior in Virtualized and Multi-Tenant Architectures

In the cloud, workloads frequently share compute, memory, and storage resources with other tenants or services. This introduces contention and unpredictable variations in:

- CPU availability,
- Buffer cache performance,
- Network bandwidth.

Such interference can cause sudden spikes in response times, affecting both batch jobs and transactional queries.

### 2.5 Legacy Optimizer Statistics and SQL Plan Regressions

Another major contributor to post-migration tuning issues is the continued reliance on legacy execution plans, hints, or optimizer statistics. These were often fine-tuned for the original on-premises environment and may no longer be optimal in the cloud. The resulting symptoms include:

- Suboptimal join methods (e.g., nested loops instead of hash joins),
- Ignored indexes or redundant full-table scans,
- Poor cardinality estimates leading to incorrect plan choices.

Unless corrected, these regressions can persist across reboots and workloads, compounding performance problems.

### 2.6 Reactive and Manual Tuning Limitations

Traditional tuning techniques, while effective in static environments, become increasingly unmanageable post-migration. Manual AWR report reviews, SQL tuning advisor suggestions, and custom scripts do not scale in cloud environments characterized by elastic workloads and constantly shifting baselines. Furthermore, the delay between problem detection and resolution increases mean time to resolution (MTTR) and creates friction in operational support.

Challenge Category	Description	Impact
Hardware Disparity	Exadata vs. commodity VM performance	Query execution delays
I/O Bottlenecks	Limited throughput, misaligned storage tier	Slow read/write latency
Cloud Config Issues	Misconfigured shapes, SGA/PGA, autoscaling	Memory contention, poor scaling

Multi-Tenant Architecture Effects	Resource sharing and contention	Unpredictable response times
SQL Plan Regression	Legacy hints/statistics causing inefficient plans	Suboptimal joins, full scans
Manual Tuning Gaps	Reactive and human-intensive tuning	Increased MTTR, scalability gap

### 3. AI-Driven Performance Tuning: Architecture and Use Cases

Artificial Intelligence (AI) has emerged as a critical enabler in overcoming the limitations of traditional performance tuning in cloud-based Oracle environments. AI enhances visibility, automates decision-making, and introduces predictive and adaptive capabilities into performance management workflows. This section presents the high-level architecture of AI-assisted tuning frameworks and explores key use cases across various layers of the Oracle technology stack.

#### 4.1 Architecture of AI-Augmented Tuning Framework

A typical AI-augmented performance tuning architecture in a post-migration Oracle cloud environment consists of the following layered components:

**Figure 1. AI-Augmented Oracle Tuning Architecture (to be generated)**

Layer	Components	Role
<b>Data Collection</b>	AWR/ASH, OS telemetry, SQL monitor, OCI logs	Captures historical and real-time performance metrics
<b>Ingestion &amp; Storage</b>	Log aggregators (OCI Logging, CloudWatch), Data Lakes	Stores and indexes logs and telemetry for analysis
<b>AI/ML Engine</b>	Oracle AI models, ML algorithms, anomaly detectors	Analyzes trends, detects anomalies, and predicts performance regressions
<b>Automation Layer</b>	Auto-indexing, SQL plan evolution, scaling policies	Implements recommendations or remediations automatically
<b>Feedback Loop</b>	Telemetry updates, human review, retraining models	Continuously refines model accuracy and tuning decisions

This layered approach allows for holistic observability and real-time adaptive optimization across compute, memory, I/O, and SQL layers.

#### 4.2 Use Case 1: Autonomous SQL Tuning

Oracle Autonomous Database leverages AI to:

- Identify slow-running SQLs,
- Recommend and implement indexes,

- Use adaptive execution plans based on run-time statistics.

AI automatically applies plan directives, evolves SQL profiles, and tracks regressions without manual intervention. DBAs can review the impact via performance hubs or Oracle Cloud console.

**Example:** In a post-migration OLTP system, AI-driven automatic indexing led to a 60% reduction in average query response time for high-frequency SELECT operations.

#### ***4.3 Use Case 2: Anomaly Detection with AWR and ASH + AI***

AI algorithms can enhance AWR/ASH data by identifying:

- Deviations in wait event patterns,
- Sudden spikes in CPU, I/O, or contention metrics,
- Rare or emergent SQL execution behaviors.

Using clustering, time-series analysis, or unsupervised learning, AI detects problems not evident through static reports.

**Example:** During a load spike in an OCI-hosted finance database, AI detected increased “log file sync” wait times caused by unbalanced commit operations, enabling preemptive load rebalancing.

#### ***4.4 Use Case 3: SQL Plan Evolution and Insight***

Tools like Oracle SQL Performance Analyzer (SPA) and Real Application Testing (RAT) can be paired with AI/ML models that:

- Learn from historic plan executions,
- Flag inefficient plans or regressions,
- Suggest plan baselines or optimizer hints.

This enhances predictability and reduces the risk of performance regression during upgrades or patching.

**Example:** In a healthcare analytics platform migrated to Exadata Cloud@Customer, AI-assisted plan regression analysis helped reduce bad plan deployment by 90%.

#### ***4.5 Use Case 4: Predictive Resource Scaling and Capacity Planning***

AI models trained on workload history can forecast resource consumption and suggest:

- New VM shapes or configurations,
- Increased IOPS or storage tier transitions,
- CPU/memory scaling policies.

In OCI, AI-based autoscaling ensures cost-efficiency by provisioning just enough resources to meet SLAs without overprovisioning.

**Example:** A retail supply chain database on OCI scaled compute nodes pre-emptively during a seasonal demand spike, maintaining sub-second response times with 30% less cost.

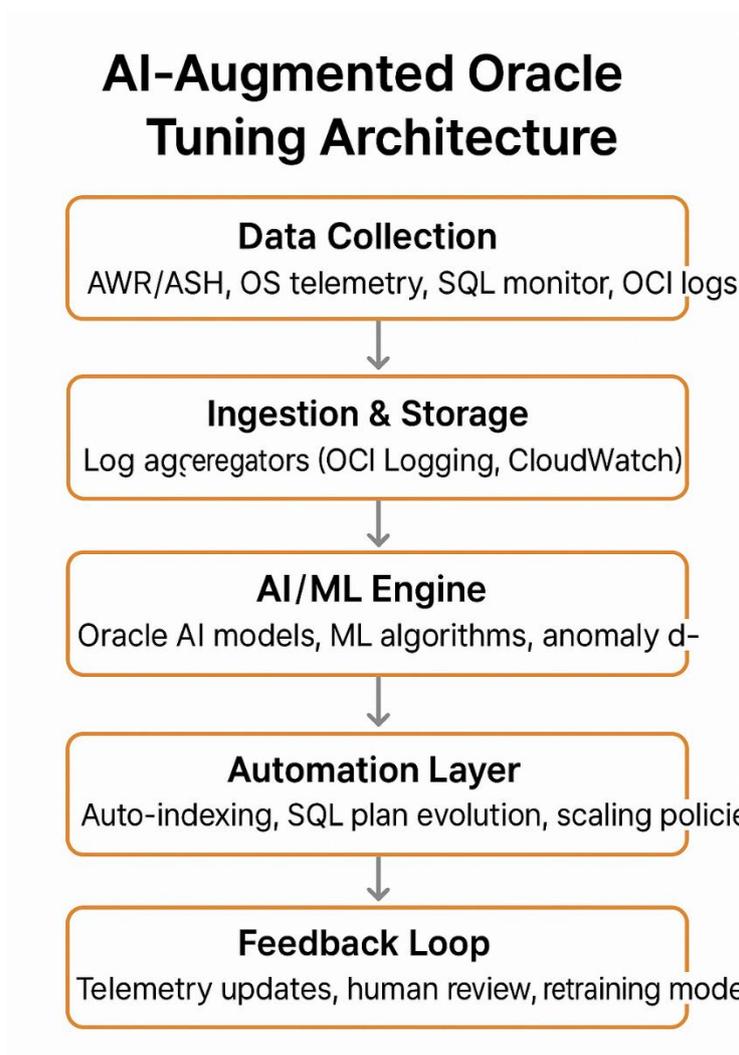
#### 4.6 Use Case 5: Automated Root Cause Analysis (RCA)

Using correlation engines, AI can link symptoms (e.g., slow query) to causes across layers:

- OS-level CPU saturation,
- Storage IOPS bottlenecks,
- Locking/blocking sessions.

Oracle Cloud Observability and Management Platform integrates AI for RCA, reducing Mean Time to Resolution (MTTR) from hours to minutes.

**Example:** In a hybrid cloud deployment, AI correlated slow INSERT operations to unexpected changes in ASM redundancy policy, which was quickly corrected.



#### 4. Comparative Analysis of Traditional vs AI-Driven Tuning Approaches

As organizations transition their Oracle workloads to the cloud, the limitations of traditional tuning techniques become increasingly evident. In contrast, AI-driven tuning introduces predictive, automated, and scalable solutions that are more aligned with the demands of elastic cloud environments. This section provides a comparative evaluation of both paradigms across critical performance dimensions.

##### 4.1 Key Dimensions of Comparison

The effectiveness of performance tuning techniques can be assessed based on the following dimensions:

- **Detection Accuracy** – Ability to identify the root cause of performance issues.
- **Response Time** – Time taken from issue identification to remediation.
- **Scalability** – Suitability for large-scale or dynamic workloads.
- **Automation Level** – Degree of manual intervention required.
- **Adaptability** – Ability to handle changing workload patterns or system configurations.
- **Proactive vs. Reactive** – Whether issues are addressed before they impact users.

##### 4.2 Comparative Summary Table

Criteria	Traditional Tuning	AI-Driven Tuning
<b>Detection Accuracy</b>	Depends on DBA expertise and static metrics	Uses ML to detect anomalies across time-series and multidimensional data
<b>Response Time</b>	High (manual analysis, remediation scripts)	Low (real-time detection, automated response)
<b>Scalability</b>	Limited in large or elastic environments	Designed for scalable, distributed cloud workloads
<b>Automation Level</b>	Minimal (scripts, manual SQL tuning)	High (auto-indexing, auto-scaling, plan evolution)
<b>Adaptability</b>	Poor adaptability to new patterns or changes	Continuously learns and adjusts to new workloads
<b>Proactive Capability</b>	Mostly reactive (post-failure diagnostics)	Proactive (predictive tuning, early alerting)

##### 5.3 Case Illustration: SQL Plan Optimization

###### Traditional Approach:

A DBA identifies a long-running query, manually examines its execution plan using EXPLAIN PLAN, compares it to previous baselines, and may apply hints or create indexes. The process can take hours to days depending on query complexity and system load.

**AI-Driven Approach:**

AI models continuously monitor query execution statistics. When a deviation from the baseline is detected (e.g., plan change leading to higher I/O cost), the system automatically recommends or reverts to a better-performing plan. Oracle Autonomous Database handles this without human intervention.

**5.4 Operational Impact**

The move from traditional to AI-enhanced tuning shifts the role of the DBA from executor to overseer. Instead of manually managing execution plans or performance alerts, the DBA reviews AI-recommended actions, focuses on architectural tuning, and handles edge cases where domain-specific knowledge is required.

**Example Impact:**

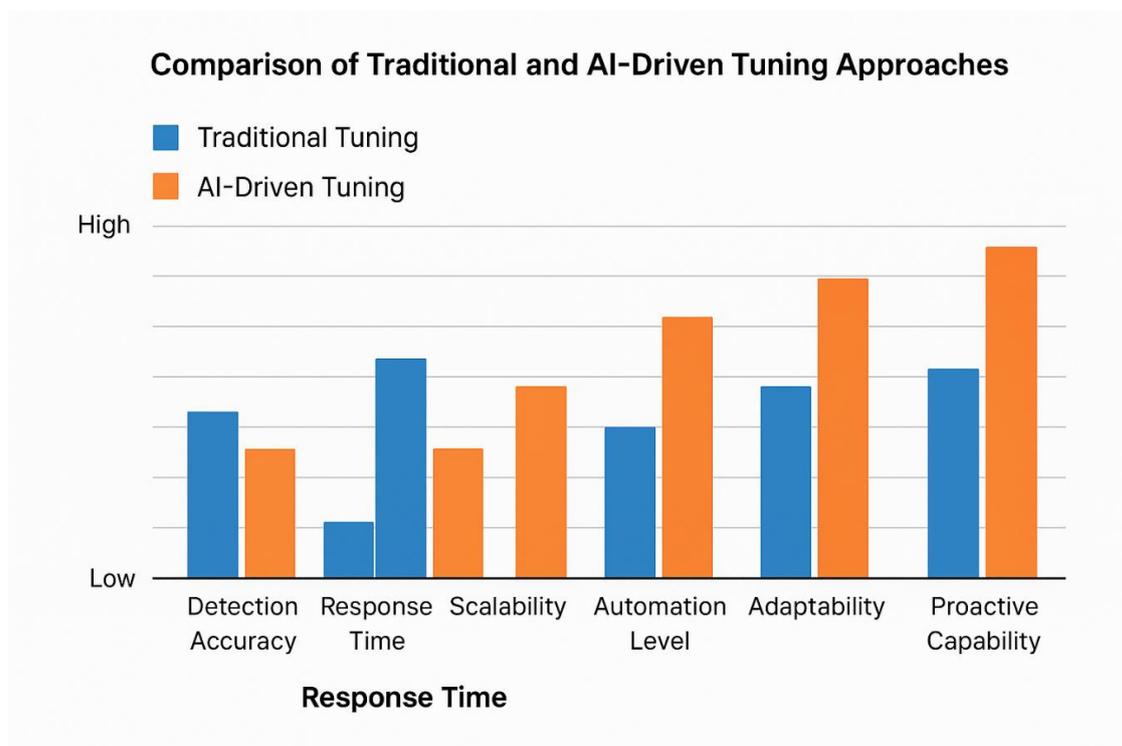
- **MTTR Reduction:** Mean Time to Resolution can drop from 4–8 hours to <30 minutes.
- **Efficiency Gains:** Up to 50–70% reduction in time spent on repetitive tuning tasks.
- **Performance Consistency:** AI tuning minimizes the likelihood of regressions post-upgrade or post-deployment.

**5.5 Challenges in AI Adoption**

Despite its benefits, AI-driven tuning has some constraints:

- **Model Explainability:** DBAs may distrust black-box recommendations without clear logic or traceability.
- **False Positives/Negatives:** Poorly trained models can misclassify performance trends.
- **Operational Readiness:** Organizations may lack maturity or observability pipelines to support AI integration effectively.

Successful adoption depends on combining AI with human oversight, establishing feedback loops, and gradually increasing automation confidence over time.



## 5. Best Practices for Oracle DBAs in AI-Enabled Environments (Condensed)

As AI becomes integral to Oracle performance management, DBAs must adapt by combining domain expertise with AI-powered automation. This section outlines practical strategies to ensure effective and responsible adoption.

### 5.1 Establish Baselines Before Migration

Capturing performance baselines pre-migration helps identify regressions post-migration. Use tools like AWR, RAT, and SQL Performance Analyzer (SPA) to document query execution times, optimizer statistics, and I/O profiles.

### 5.2 Enable AI Features Early

Activate AI capabilities such as automatic indexing and autonomous SQL tuning during or immediately after migration. Tools like Oracle Autonomous Database and Oracle Management Cloud offer early insight into workload behavior.

### 5.3 Review AI Recommendations Before Applying

Not all AI suggestions should be applied blindly. For critical SQL statements, validate proposed changes using SQL Tuning Advisor or test environments. Maintain manual oversight to ensure accuracy and predictability.

#### ***5.4 Use Feedback Loops to Improve AI Accuracy***

Regularly review outcomes of AI actions and tag successful or ineffective recommendations. This helps refine models and ensures that automation adapts to evolving workloads.

#### ***5.5 Focus on Strategic Tasks***

Allow AI to handle repetitive tuning while DBAs focus on architecture decisions, cost optimization, and performance planning. Upskill in areas like observability, AIOps tooling, and workload pattern analysis.

### **7. Conclusion**

The migration of Oracle workloads to cloud platforms introduces a range of performance challenges that cannot be effectively addressed through traditional tuning methods alone. Variations in infrastructure, workload behavior, and platform configurations create a complex tuning landscape that requires more than manual diagnostics and reactive interventions.

Artificial Intelligence (AI) has emerged as a transformative solution in this space, enabling autonomous tuning, intelligent anomaly detection, predictive resource scaling, and automated root cause analysis. By integrating AI into Oracle performance management—via tools such as Oracle Autonomous Database, Oracle Management Cloud, and third-party AIOps platforms—organizations can not only reduce Mean Time to Resolution (MTTR) but also ensure consistent and scalable performance across dynamic, cloud-native environments.

Importantly, AI is not a replacement for Oracle DBAs, but a force multiplier. It enables DBAs to shift their focus from low-level troubleshooting to high-value strategic tasks such as architecture optimization, cross-stack observability, and long-term performance planning. By adopting AI thoughtfully and aligning it with best practices, organizations can unlock the full potential of their Oracle cloud investments while future-proofing their operational models.

### **8. References**

Here is a starter list of references. You may add more based on journal formatting (APA/IEEE/Elsevier/etc.) and cited tools or datasets:

- [1] Oracle Corporation. (2023). *Oracle Autonomous Database Technical Overview*. Retrieved from <https://www.oracle.com/autonomous-database/>
- [2] Oracle Cloud Infrastructure. (2023). *OCI Observability and Management Platform*. Retrieved from <https://docs.oracle.com/en-us/iaas/management/>
- [3] Dynatrace. (2023). *AI-Powered Performance Monitoring for Oracle*. Retrieved from <https://www.dynatrace.com/>
- [4] AppDynamics. (2023). *Database Visibility for Oracle Workloads*. Retrieved from <https://www.appdynamics.com/>
- [5] Datadog. (2023). *Monitoring Oracle Databases in Hybrid Environments*. Retrieved from <https://www.datadoghq.com/>

**Citation:** Prasad Manda. (2023). Leveraging AI to Improve Performance Tuning in Post-Migration Oracle Cloud Environments. *International Journal of Artificial Intelligence & Machine Learning (IJAIML)*, 2(1), 235-247.

**Abstract Link:** [https://iaeme.com/Home/article\\_id/IJAIML\\_02\\_01\\_021](https://iaeme.com/Home/article_id/IJAIML_02_01_021)

**Article Link:**

[https://iaeme.com/MasterAdmin/Journal\\_uploads/IJAIML/VOLUME\\_2\\_ISSUE\\_1/IJAIML\\_02\\_01\\_021.pdf](https://iaeme.com/MasterAdmin/Journal_uploads/IJAIML/VOLUME_2_ISSUE_1/IJAIML_02_01_021.pdf)

**Copyright:** © 2023 Authors. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Creative Commons license:** Creative Commons license: CC BY 4.0



✉ [editor@iaeme.com](mailto:editor@iaeme.com)