

Optimizing Cloud Computing Performance With Advanced DBMS Techniques: A Comparative Study

Janardhana Rao Sunkara^{1*}, Sanjay Ramdas Bauskar², Chandrakanth Rao Madhavaram³,
Eswar Prasad Galla⁴, Hemanth Kumar Gollangi⁵

^{1*}CVS Pharm. Inc. Sr. Oracle DB Admin, JanardhanaRaoSunkara@outlook.com

²Pharmavite LLC Sr. Database Administrator, sanjayramdasbauskar@outlook.com

³Microsoft Support Escalation Engineer, Craoma101@outlook.com

⁴Dept. of Comp. Sci. Univ. of Central Missouri, EswarPrasadGalla@outlook.com

⁵Dept. of Comp. Sci. Univ. of SouthEast Missouri, HemanthKumarGollangi12@outlook.com

Abstract

In the era of digital transformation, optimizing cloud computing performance has become a critical focus for organizations striving to leverage the full potential of cloud infrastructures. This study presents a comparative analysis of advanced database management system (DBMS) techniques aimed at enhancing cloud computing performance. By examining a range of strategies, including indexing optimizations, query performance tuning, data partitioning, and caching mechanisms, the research identifies key methodologies that can significantly impact efficiency and scalability in cloud environments. Through a series of tests and performance metrics, this study evaluates the effectiveness of these techniques across various cloud platforms and workloads. The findings provide valuable insights into which DBMS approaches offer the greatest benefits in terms of speed, resource utilization, and overall system performance. This comparative study not only highlights the strengths and weaknesses of different techniques but also offers practical recommendations for organizations seeking to optimize their cloud computing infrastructure.

Keywords: Cloud Computing Performance, Advanced DBMS Techniques, Database Management Systems, Performance Optimization, Cloud Database Optimization, Comparative Study, DBMS Performance Tuning, Cloud Resource Management, Scalability in Cloud Computing, Distributed Databases, Query Optimization, Data Storage Solutions, Cloud Infrastructure, Database Scalability, Data Access Efficiency, Cloud Performance Metrics, Advanced Database Techniques, Load Balancing in Cloud, Database Query Processing, Cloud-Based Database Management, Data Management Strategies, Performance Benchmarking, Cloud Application Performance, DBMS Efficiency, Cloud Storage Optimization.

1. Introduction

The primary goal of this study is to empirically ascertain whether cloud computing performance can be improved through practical methods in the database management system (DBMS) that employ basic and classic indexing techniques. Experimentation on complex real-world queries is performed during the project in order to validate these techniques. DBMS are cloud-unaware, resulting in inefficient query optimization. In this atmosphere, the surrogate-based technique performs effectively, but substantial knowledge of the data is necessary. Basic translation and indexing methods, on the other hand, do not necessitate substantial knowledge of the saved data. Even with large amounts of data, advanced techniques provide results similar to basic embedding and indexing techniques.

In modern times, monetization and digitalization have emerged as new trends, emphasizing the fundamental importance of data in this context. According to statistics, the volume of data generated each day in the world is equivalent to 2.5 quintillion bytes, and 90% of the global data to date has been generated within the last two years.

As a result, these trends have shifted much of the computing load to the cloud. Because cloud computing is a multi-tenant environment with non-shareable computing resources, database management system (DBMS) throughput and response times are important performance indicators. Despite making progress over the past several years, DBMS still faces a number of challenges when run in cloud resources.




Fig 1 : Cloud Computing and Big Data

1.1. Background and Rationale

The cloud computing paradigm promises to provide a flexible IT infrastructure like hardware, software, and other resources but governed by strict Service Level Agreements (SLAs). However, the performance of cloud applications can be highly affected by the performance of the underlying Database Management Systems (DBMS). Improving the DBMS performance will lead to a subsequent increase in the cloud application performance. The authors tried to improve the performance by increasing the number of research nodes and storing the frequently accessed data locally, in order to decrease the remote read operations. Other researchers have studied and analyzed the performance of data placement strategies in the cloud. Other studies have been devoted to modeling and predicting the end-to-end performance of cloud/DBMS applications based on cluster resources such as CPU, I/O, and so forth. An important research direction that has not been completely investigated in the literature is how the architecture of the underlying DBMS, such as main memory, backup storage, indices, indexing structure, and recovery technique, can affect cloud computing performance. Such DBMS features and all the architectural decisions can have a significant impact on the amount of appropriate main memory assigned for processing the transactions as well as serving the ad-hoc queries. Since the DBMS should recover all the data accessed/used by the end user, backup storage architecture decisions are also important. In this paper, our target is to study and compare different indexing and backup/recovery techniques considering their impact on the overall system performance.

Equ 1: Multi-objective fitness function ($J \rightarrow (\cdot)$)

Input : Search agent, ϵ_i

Output: Criterion function, $\vec{J}(\mathcal{X}_i) = \{J_1(\mathcal{X}_i) \quad J_2(\mathcal{X}_i)\}$

```

1 Assign  $i^{th}$  solution to null vector, i.e.,  $\mathcal{X}_i \leftarrow \varnothing$  and  $\xi_i \leftarrow 0$ 
  */      Decode parent
2 for  $m = 1$  to  $n$  do
3   if  $\epsilon_{i,m} = 1$  then
4     Computed fitness function using Eq. 3.
5      $\mathcal{X}_i \leftarrow \{\mathcal{X}_i \cup x_m\}$ 
6      $\xi_i \leftarrow \xi_i + 1$ 
7   end
8 end
9 Compute coefficients ( $\Theta$ ), corresponding to  $\mathcal{X}_i$ 
  */      Evaluate the criterion function
10 Determine the dynamic prediction error ( $\mathcal{E}_i$ )
11 Determine the penalty assignment ( $\mathcal{P}$ )
12  $J_1(\mathcal{X}_i) \leftarrow \xi_i + \mathcal{P}$ ,  $J_2(\mathcal{X}_i) \leftarrow \mathcal{E}_i + \mathcal{P}$ 
```

2. Cloud Computing and DBMS: An Overview

In the modern world, several businesses are rapidly shifting toward cloud computing to reduce costs, increase operative dexterity, and distribute office efficiency. Cloud computing possesses several eye-catching advantages: flexibility,

security, cost, consistency, and tons of others. There are several contrasts between on-site and cloud imperatives. In the underlined part, one of the requisite contrasts oversees apparent on-cloud imperatives. Cloud is an environment with so many stakeholders providing a wide variety of applications. The client/customer connects to the application that he/she is interested in by means of a cloud account. One of the most common applications of cloud computing is databases which are managed as a service for the client and the operations are executed for them in the cloud environment. Hence, with the increasing number of users, databases are becoming more and more important. The advancement of cloud computing combines a database management system (DBMS) into a pay-as-you-go idea in which the DBMS consumers can elude the burden of purchasing hardware and software. DBMS bequest has an appositive effect on cloud computing domains. Database as a service (DBaaS) is one of the pre-eminent, high-intensity stuff of cloud computing. The lucrative sides of DBaaS are the replacement of capital charges with operating charges, better-quality output, separate overhead, and the power to deploy arduous operations at a faster tempo.

Operating conventional DBMS in cloud nodes purely does not insinuate optimal conduct due to several doings and connections in the cloud globe. This encyclopedia paper submits an educational road to optimizing DB performance in cloud facilities (IaaS). The method helps a decision maker (DBA) to choose a cloud DBMS using the optimal DB selection method. The publications used in this paper were found using the systematic research method. The paper is organized as follows. Section 2 briefly depicts cloud computing and SaaS as well as their significance, allegiance to solving IT inelegance, and breeze DBSaaS models. Allocating regulations from an economic-geographic view between stakeholders and the suspected order delayed cloudy environments and adventitious pressures of global justice issues on multifold tiers. Privacy policies, the aptness of legislation to challenge jurisdictional conflicts, international law, multicultural, multilingual, complex delegation, technology threats, intentional disorientation, and other issues make cloud computing laws multifaceted. Since the cloud operating paradigms' long-existing and multifaceted regulations are non-stop from the legal surveillance and fruitful network cloud locus. In the contemporary business landscape, cloud computing has become a pivotal force driving efficiency and cost-effectiveness. By leveraging the cloud, companies can enjoy a range of benefits, including flexibility, enhanced security, and consistent performance. One prominent application within this realm is Database as a Service (DBaaS), which allows organizations to utilize database management systems (DBMS) on a pay-as-you-go basis, alleviating the need for upfront hardware and software investments. This shift from capital expenditures to operating expenses, coupled with the ability to scale operations rapidly, positions DBaaS as a significant advantage in cloud computing. However, operating traditional DBMS in cloud environments presents challenges due to the complexities of cloud infrastructure. To address these issues, an educational approach to optimizing DB performance in Infrastructure as a Service (IaaS) environments is crucial. This approach aids decision-makers in selecting the most suitable cloud DBMS by leveraging systematic research methods. The discussion extends to the broader implications of cloud computing laws, which encompass privacy policies, jurisdictional conflicts, and international regulations, reflecting the intricate nature of legal compliance in a globalized cloud landscape.




Fig 2 : Overview of Cloud Computing

2.1. Key Concepts

Cloud computing delivers storage, network, and impactful computation services over the internet. It brings substantial benefits such as high fault tolerance and resources as services. Services can be recovered if a cloud provider at any time takes back resources. In contrast, the database management system (DBMS) can be defined as a database person's copilot. At the primary level, it is a software application. At the secondary level, it talks to user applications and the database. However, the DBMS efforts to provide performance security are defined as the concept of ensuring that all important factors of database systems work to equip the database system with consistent reliability, including serving large numbers of users. An upgrading system can be nonetheless good. There are many other concepts associated with the DBMS, but authentication is critical. Efficiency and inefficiency are utilized best for strategically designing the database shape and keeping the database pieces that are posterior.

Cloud computing provides storage, networking, and computing services; a cost-efficient, consumer-centric, pay-as-you-go, on-demand, web-based utility model is on the internet. Resources: SaaS, PaaS, and IaaS provide three categories (security, privacy, legal, and reliability) of cloud applications. The data uses other different service applications. The provision of cloud database services has become increasingly popular at present. Cloud databases are offered as services that need little or no hardware or software — the application rather than the cloud database operates. The database management system (DBMS) is a very large problem for use in hardware and operating systems. Database for the physical and logical data is to securely store and retrieve effectively. The latest server-based and scale-out databases cloud workloads that developed architectures have allowed storage to add cost-effectively and ensure public/hybrid and multi-cloud data management and portability. Repositories, data warehouses, and data marts are used in conjunction with a wide industry running a larger cloud database. Baseline manufacturers, until a decade ago, primitive data by conventional file systems, mix-and-match, or are widely used in a single computational cloud is a standard database. The cloud computing climate characteristics and common DBMS technologies and other use cases would include resilience, suboptimal performance, and failure to adapt to storage tuning practices. To handle more cloud workloads, their algorithms running the GPU have been updated.




Fig : Key concepts for measuring elasticity

3. Performance Optimization in Cloud Computing

One of the primary advantages offered by cloud computing includes virtually limitless resources, be they computational, data storage, or even software and applications. In the contemporary IT landscape, cloud computing is a technological reality, and big data analytics is its killer application. However, the presence of a plethora of services exponentially deteriorates the scope of security and reduces the performance through resource contention, which is a serious issue in IaaS and PaaS models. Performance optimization in the cloud environment is arduous due to the plethora of potential combinations of candidates involving virtual machines and data management systems, e.g., database management systems (DBMS) among others. The two natural paradigms of enhancing cloud performance encompass the optimization of data management systems (DMS) and the efficient management of resources. Efficient management of database management systems (DBMSs) is a precursor of reducing the query processing overhead and correlates as a non-negligible aspect of enhancing cloud performance. In view of this, researchers and industry professionals have obtained a paradigm shift from conventional database management system features such as concurrency control, recoverability, and error-free processing of workloads in the presence of abilities, e.g., white-box and black-box optimizers among others, that pave the way for

developing an optimized query plan with the least processing latency. Therefore, by mitigating such overhead, the final objective of reducing the processing time of the workloads may be achieved in a meaningful manner. The DBMS internal enhancements involve indexing, materialized views, caching, query optimization, etc. However, resource brokers or managers heavily rely on the optimal assignment of virtual machines, which results in less resource contention and precludes resource monopolization.




Fig 3 :Cloud computing-based optimization

3.1. Challenges and Solutions

There are many challenges to optimizing the performance of cloud computing. For example, traditional database management systems are not designed to be utilized in cloud computing environments, and challenges arise when the processing load is shared across multiple distributed and commodity hardware systems. Performance bottlenecks and interdependencies between database and application workloads make monitoring and tuning complex. In addition, with the increasing number of server nodes in cloud environments, a time loop is needed to schedule the physical resources required for processing database execution tasks and delivering the result considering multiple logical channels via the network hardware. Two ways to eliminate or mitigate these challenges. First, solutions break down the performance tuning overhead into multiple time-distributed feedback controls, each with a specific performance optimization scope. The timing control hierarchy concurrently tracks the performance of the workload and relational database management system (DBMS) at the service instance, virtual machine, and host system levels. Centralized data collection provides the diagnostic feedback variables. The outputs of the feedback controllers are control signals that direct multiple control subprocessors, each of which applies the performance optimization techniques described in this paper, thereby providing dynamic adaptability to the changing performance demands placed on the cloud resources. Control loop cross-coupling information is used to modify the adaptive algorithms so that each independent feedback control response will impact the future performance subsystem response in a manner that collaboratively serves the multi-tier e-commerce system performance as a whole to leverage the increasing number of server mules.

Equ 2: Cost-Benefit Analysis Formula

Cost Benefit Analysis Formula

$$\text{Net Present Value (NPV)} = \sum \text{Present Value of Future Benefits} - \sum \text{Present Value of Future Costs}$$

$$\text{Benefit Cost Ratio} = \frac{\sum \text{Present Value of Future Benefits}}{\sum \text{Present Value of Future Costs}}$$

4. Advanced DBMS Techniques

The sheer growth of data in digitized environments has ushered the database management system industry to enter into cutting-edge areas. This necessitated various advanced techniques for the creation and management of databases, with a focus on the need for agility and high performance. Furthermore, the availability of advanced hardware and network

technologies facilitated the deployment of databases on a large scale. Cloud computing technologies have also gained widespread popularity because of their low costs and scalability, leading to the sharpening of the focus on DBMS techniques relative to cloud computing environments to reduce response time. Some of the characteristics of database systems such as ACID properties, response time, etc., are seriously affected when transaction requests are very large or are mingled with analytic workloads. The advancement in disk drives, but fewer improvements in data access time, initiated the quest for in-memory databases as a potential solution. Another development in database management systems came as a result of the growth activities in terms of IoT and Big Data. It appears that it is important to analyze large datasets before making a decision. As a result, research communities have been working on new designs for cloud computing databases. It seems that relational databases do not suit the concept fully in this case, leading to the commencement of the analysis of analytic databases. The aforementioned databases are specifically designed based on specific hardware settings and are crucial for extracting information. There is a strong need to conduct a comparative study to see the effects of these databases from a cloud perspective. This study analyzes two types of databases, one of which is an in-memory database with traditional row-level storage techniques, and the other is an analytic database, which is columnar-based.




Fig 4 : Database management system performance comparisons

4.1. In-Memory Databases

One of the main techniques used for improving database performance is utilizing main memory for data storage and retrieval, which is common in in-memory databases. The concept of in-memory databases is not new, but in the recent past, due to advancements in both hardware and software, in-memory databases have been attracting both the research community and the commercial database vendors. Because cloud computing has become an attractive computing model for both industry and academia, researchers have started looking at using in-memory databases to bring massive benefits to large analytics workloads. However, as of now, there is very little information on how in-memory databases perform and scale in the cloud environment, among the big three cloud providers, i.e., Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP). To make things worse, in-memory engines are different, with no clear guidelines on which to use or when to use them. Hence, there is a clear need to understand how in-memory databases perform at scale in the cloud environment for large analytics workloads using multi-database transactional benchmarks. Given these challenges in a managed cloud environment, the need for in-depth evaluation of in-memory databases in the cloud to assess their potential has attracted widespread attention. With the potential, our goal is to compare different in-memory databases from AWS, Azure, and GCP using TPC benchmarks.

4.2. Columnar Databases

Traditional DBMS storage is organized by rows. This paper primarily investigates a new concept focused on columnar databases, which have been defined as a recent column-oriented storage technique in which column families are stored vertically and consecutively, such that all values of column 1 are stored together, all values of column 2 are stored together and so on. Moving from row-wise storage to columnar-wise storage is proposed to bring improvements in the scan time from seconds to milliseconds, leading to the performance gained in cloud computing. Increasing the performance of such queries is expected to improve the overall performance of a range of applications. In cloud computing, cloud-based data management combines two emerging technologies: cluster computing architectures, typically available in a cost-competitive on-the-fly manner, and complex data management systems. The process of storing, grooming, visualizing, and enabling the HIPO concert analysis by employing column-based processing techniques on massively-scale histograms

will potentially produce an order of magnitude savings in time and data transfer costs. This setup is especially useful for the case where analysis happens in cloud storage, and therefore minimizing memory/input/output requirements becomes critical. The inefficient use of databases can bottleneck the I/O cycle. The typical strategy of reclaiming space in database management systems (DBMS) is to replicate the table, perform transformations on the copy, and replace the original tables once permanently entangled and consistent with the query update propagation.




Fig : Impact of the number of nodes on performance

5. Methodology

In this section, we present the research design and approach of the paper. The overall research methodology proposed in this paper aims to empirically prove the research hypothesis introduced earlier. Concretely, research in this paper is conducted through an experimental approach by conducting an in-depth performance evaluation of IBM DashDB (Giants) as the technology representative of the advanced DBMS in comparison with the traditional DBMS. There are several methodologies used in this study to perform the performance comparison of massively parallel DBMS in a cloud computing environment. First, a cost model aggregation of the three following factors was computed. The service cost computation was done at the true-up usage period. Second, the performance comparison was performed in an on-demand situation by concurrently running complex SQL workloads to assess the performance in terms of response time. Thirdly, to verify the accuracy of performance comparison results, the complexity of various parts of the advanced DBMS and traditional DBMS will be compared. The response time performance comparison of the two different technologies of DBMS is to be performed using three phases. The first experiment is the experimental database schema design, where we will consider in which areas the DB technologies differ from one another. The second experiment will focus on validation. It is about checking which parts of the advanced DBMS are faster than others, so it will not be affected by all DB microservices. Moreover, another aspect that can be tested is the concurrency. For example, in the testing situation for the throughput, it is possible to configure one service to use one faster DB functionality module. The last performance comparing the experiment will be based on the experimental methodology in the cloud. The service-made cost represents the metrics of the DBMS workload intensity and duration. Therefore, we consider both factors in the performance testing. DBMSs also introduced the concept of "elastic service", where the extent of one or several computing resources is limited by using less and less of the rest of them. The research design of this paper is structured to empirically validate the hypothesis regarding the performance capabilities of advanced DBMS, specifically IBM DashDB (Giants), in contrast to traditional DBMS within a cloud computing framework. This study employs an experimental methodology encompassing a comprehensive performance evaluation across several dimensions. Initially, a cost model aggregation is established, considering service costs during the true-up usage period. The performance assessment is conducted in real-time, utilizing complex SQL workloads to gauge response times under concurrent execution. To ensure the robustness of the findings, the study meticulously examines the intricacies of both advanced and traditional DBMS, focusing on specific components to identify performance differentials. The experiments unfold in three phases: first, through the design of an experimental database schema to delineate the distinctions between the DB technologies; second, by validating the performance of various advanced DBMS components in isolation to mitigate microservice interference; and finally, by evaluating performance within a cloud environment, taking into account the metrics of workload intensity and duration. This multifaceted approach underscores the innovative concept of "elastic service" introduced by modern DBMS, emphasizing efficient resource utilization while maintaining performance integrity.

Equ 3: Optimal load balancing and assessment of existing load balancing criteria

```

Input: numIter: the number of iterations to compute, priorityQueue: a
        priority queue
1 foundLb[i] = 0  $\forall i = 1..K$ ;
  // root node
2 cNode = Node(iter=0, LB=true, cost=0.0, appState, lbState, prev= $\emptyset$ );
3 while cNode.iter < numIter do
4   if cNode.LB then
5     | foundLb[cNode.iter] = true;
6   end
7   dontLbNode, doLbNode = cNode.getChildren();
8   if not foundLb[doLbNode.iter] then
9     | // Measurement of cost (i.e., time) with a
10    |   theoretical model or a real application
11    | doLbNode.computeCost();
12    | replaceOrInsertNode(priorityQueue, doLbNode);
13  end
14  dontLbNode.computeCost();
15  insert(priorityQueue, dontLbNode);
16  cNode = priorityQueue.pop();
17 end

```

5.1. Research Design and Approach

Our comparative study is guided by the need to ensure an optimal configuration of the DBMS for the purpose of reducing the execution time required to process the expected workloads. Within this context, we perform the correction of a proposed DBMS configuration and assess the implications in terms of performance. The study is based on an experimental assessment using the cloud-specific DBMS engine, Amazon Aurora. Table 5.2 highlights the research design and approach containing data collection, data analysis, and comparison procedures.

- Data collection - in this phase, the experimental environment and the setup from prior similar works were used in the data collection. Given that cloud-based technologies were used in our study, we selected five TPC-C workloads, which are considered suitable for cloud deployment purposes. The data will be collected from the implemented workload performances with different available resources and different numbers of VCPUs.
- Data analysis - for data analysis, a number of 50 transaction performance results were published (in milliseconds) from different workloads and then the results were compared by assessing the 95% confidence intervals of the means for 95% confidence in normal distribution. The outcome of the research, as a result of assessing the 95% confidence interval, shows that the performance difference between two means will be considered statistically significantly different if there are no overlapping confidence intervals.
- Step 4 - Data comparison: The statistical experiment versus the TPC-C performance results is conducted to compare the proposed advanced techniques. These techniques were the impact of using the third normal form against column-based technology. In addition, an advanced database system feature was exploited for the purpose of enhancing faster SQL execution, which is querying the column store indexes (via In-Memory OLTP) in parallel.

6. Conclusion

This paper conducts a comparative study to evaluate the performance of the two widely used advanced DBMS techniques. To this end, we use block nested loop join using symmetric key operations in SQL Server and Intel SGX supported in the SQL Server database engine. Extensive experimental results provide several findings. We observe that SQL Server takes less time when the sizes of data selection are 1% and 10%. Intel SGX takes less time when the size of data selection is 100%. In the presence of range query, SQL Server takes less time when the product of the sizes of data selections of related tables is greater than or equal to 1% and the size of the smaller table is small. The data sizes are assumed equivalent and performance differences are always on the order of a few seconds. More interesting findings include the fact that Intel SGX encrypted U-SQL queries perform better relative to T-SQL than unencrypted U-SQL. In this study, we use the block nested loop join algorithm and propose that all block nested loop join techniques in the real system use symmetric key operation and Intel SGX technology. To the best of our knowledge, there is no existing study that has this focus. We conduct a comparative study of the DBMS taking into account the conditional range queries and employing the SQL Server 2019. Extensive experimental results from the Azure cloud environment provide several empirical findings that can be useful for cloud developers and practitioners in the optimization of cloud computing performance using current technologies. While SQL Server is the commonly utilized DBMS, we believe that modern DBMS have comparable performance. It is always possible to compare more DBMS and new trends in DBMS. We also propose profitable future trends based on the current energy efficiency for future work.

6.1. Future Trends

The domain of query processing in the upcoming clouds will continue to grow, and they will be more sophisticated and complex. There will be a greater necessity for spatial support due to the advent of mobile devices. RDBMS firms are continuously modifying and improving their techniques in an attempt to boost their position in this largely niche sector, concentrating on performance optimization, scalability, and flexibility. Besides this roadmap presented, this paper could hopefully persuade RDBMS vendors to rethink their strategy and concept concerning the introduction of new methods in order to advance RDBMS performance. Today's optimization systems are not practical when it comes to such a huge number of candidates. In addition to several application areas, cloud computing is evolving into a full-fledged infrastructure that is anticipated to provide everything as a service, from basic processing and storage to high-level services like databases. Several cloud database systems and management services have been developed. Despite ongoing research in database management systems, many areas need further exploration. One of the new emerging issues in the field of cloud computing is exploring new techniques to optimize cloud computing performance. Database management systems (DBMS) were set to continue the development and exploration of innovative methods to enhance them. In addition, future architecture enhancements will be required to allow the incorporation of new methods and concepts.

7. References

- [1] Kommisetty, P. D. N. K. (2022). Leading the Future: Big Data Solutions, Cloud Migration, and AI-Driven Decision-Making in Modern Enterprises. *Educational Administration: Theory and Practice*, 28(03), 352-364.
- [2] Yadav, P. S. (2023). Enhancing Software Testing with AI: Integrating JUnit and Machine Learning Techniques. *North American Journal of Engineering Research*, 4(1).
- [3] Mahida, A. Explainable Generative Models in FinCrime. *J Artif Intell Mach Learn & Data Sci* 2023, 1(2), 205-208.
- [4] Vaka, D. K. (2023). Achieving Digital Excellence In Supply Chain Through Advanced Technologies. *Educational Administration: Theory and Practice*, 29(4), 680-688.
- [5] Pamulaparti Venkata, S., & Avacharmal, R. (2023). Leveraging Interpretable Machine Learning for Granular Risk Stratification in Hospital Readmission: Unveiling Actionable Insights from Electronic Health Records. *Hong Kong Journal of AI and Medicine*, 3(1), 58-84.
- [6] Chintale, P., Khanna, A., Korada, L., Desaboyina, G., & Nerella, H. AI-Enhanced Cybersecurity Measures for Protecting Financial Assets.
- [7] Avacharmal, R., Pamulaparti Venkata, S., & Gudala, L. (2023). Unveiling the Pandora's Box: A Multifaceted Exploration of Ethical Considerations in Generative AI for Financial Services and Healthcare. *Hong Kong Journal of AI and Medicine*, 3(1), 84-99.
- [8] Yadav, P. S. REAL-TIME INSIGHTS IN DISTRIBUTED SYSTEMS: ADVANCED OBSERVABILITY TECHNIQUES FOR CLOUD-NATIVE ENTERPRISE ARCHITECTURES.
- [9] Mahida, A. (2023). Enhancing Observability in Distributed Systems-A Comprehensive Review. *Journal of Mathematical & Computer Applications*. SRC/JMCA-166. DOI: doi. org/10.47363/JMCA/2023 (2), 135, 2-4.
- [10] Vaka, D. K. Empowering Food and Beverage Businesses with S/4HANA: Addressing Challenges Effectively. *J Artif Intell Mach Learn & Data Sci* 2023, 1(2), 376-381.
- [11] Pamulaparti Venkata, S. (2023). Optimizing Resource Allocation For Value-Based Care (VBC) Implementation: A Multifaceted Approach To Mitigate Staffing And Technological Impediments Towards Delivering High-Quality, Cost-Effective Healthcare. *Australian Journal of Machine Learning Research & Applications*, 3(2), 304-330.
- [12] Chintale, P., Deshmukh, H., & Desaboyina, G. Ensuring regulatory compliance for remote financial operations in the COVID-19 ERA.
- [13] Avacharmal, R., Sadhu, A. K. R., & Bojja, S. G. R. (2023). Forging Interdisciplinary Pathways: A Comprehensive Exploration of Cross-Disciplinary Approaches to Bolstering Artificial Intelligence Robustness and Reliability. *Journal of AI-Assisted Scientific Discovery*, 3(2), 364-370.
- [14] Yadav, P. S. Optimizing Data Stream Processing Pipelines: Using In-Memory DB and Change Data Capture for Low-Latency Enrichment.
- [15] Mahida, A. (2023). Machine Learning for Predictive Observability-A Study Paper. *Journal of Artificial Intelligence & Cloud Computing*. SRC/JAICC-252. DOI: doi. org/10.47363/JAICC/2023 (2), 235, 2-3.
- [16] Pamulaparti Venkata, S., Reddy, S. G., & Singh, S. (2023). Leveraging Technological Advancements to Optimize Healthcare Delivery: A Comprehensive Analysis of Value-Based Care, Patient-Centered Engagement, and Personalized Medicine Strategies. *Journal of AI-Assisted Scientific Discovery*, 3(2), 371-378.
- [17] Vaka, D. K. "Artificial intelligence enabled Demand Sensing: Enhancing Supply Chain Responsiveness.
- [18] Chintale, P., Korada, L., WA, L., Mahida, A., Ranjan, P., & Desaboyina, G. RISK MANAGEMENT STRATEGIES FOR CLOUD-NATIVE FINTECH APPLICATIONS DURING THE PANDEMIC.

- [19] Avacharmal, R., Gudala, L., & Venkataramanan, S. (2023). Navigating The Labyrinth: A Comprehensive Review Of Emerging Artificial Intelligence Technologies, Ethical Considerations, And Global Governance Models In The Pursuit Of Trustworthy AI. *Australian Journal of Machine Learning Research & Applications*, 3(2), 331-347.
- [20] Yadav, P. S. (2022). Enhancing Real-Time Data Communication and Security in Connected Vehicles Using MQTT Protocol. *Journal of Artificial Intelligence & Cloud Computing*. SRC/JAICC-E122. DOI: doi.org/10.47363/JAICC/2022 (1) E122 J Arti Inte & Cloud Comp, 1(3), 2-6.
- [21] Vaka, D. K. " Integrated Excellence: PM-EWM Integration Solution for S/4HANA 2020/2021.
- [22] Mahida, A. (2022). Comprehensive Review on Optimizing Resource Allocation in Cloud Computing for Cost Efficiency. *Journal of Artificial Intelligence & Cloud Computing*. SRC/JAICC-249. DOI: doi.org/10.47363/JAICC/2022 (1), 232, 2-4.
- [23] Tilala, M., Pamulaparti Venkata, S., Chawda, A. D., & Benke, A. P. Explore the Technologies and Architectures Enabling Real-Time Data Processing within Healthcare Data Lakes, and How They Facilitate Immediate Clinical Decision-Making and Patient Care Interventions. *European Chemical Bulletin*, 11, 4537-4542.
- [24] Chintale, P., & Desaboyina, G. (2018). FLUX: AUTOMATING CLUSTER STATE MANAGEMENT AND UPDATES THROUGH GITOPS IN KUBERNETES. *International Journal of Innovation Studies*, 2(2).
- [25] Vaka, D. K. (2020). Navigating Uncertainty: The Power of ‘Just in Time SAP for Supply Chain Dynamics. *Journal of Technological Innovations*, 1(2).
- [26] Avacharmal, R. (2022). ADVANCES IN UNSUPERVISED LEARNING TECHNIQUES FOR ANOMALY DETECTION AND FRAUD IDENTIFICATION IN FINANCIAL TRANSACTIONS. *NeuroQuantology*, 20(5), 5570.
- [27] Yadav, P. S. (2022). Automation of Digital Certificate Lifecycle: Improving Efficiency and Security in IT Systems. In *Journal of Mathematical & Computer Applications* (pp. 1–4). Scientific Research and Community Ltd. [https://doi.org/10.47363/jmca/2023\(2\)e107](https://doi.org/10.47363/jmca/2023(2)e107)
- [28] Mahida, A. Predictive Incident Management Using Machine Learning.
- [29] Pamulaparti Venkata, S. (2022). Unlocking the Adherence Imperative: A Unified Data Engineering Framework Leveraging Patient-Centric Ontologies for Personalized Healthcare Delivery and Enhanced Provider-Patient Loyalty. *Distributed Learning and Broad Applications in Scientific Research*, 8, 46-73.
- [30] Perumal, A. P., & Chintale, P. Improving operational efficiency and productivity through the fusion of DevOps and SRE practices in multi-cloud operations.
- [31] Avacharmal, R., & Pamulaparthivenkata, S. (2022). Enhancing Algorithmic Efficacy: A Comprehensive Exploration of Machine Learning Model Lifecycle Management from Inception to Operationalization. *Distributed Learning and Broad Applications in Scientific Research*, 8, 29-45.