

From Data Warehouses to Data Meshes: A Comparative Analysis of Evolving Enterprise Architectures

Srinivasa Kalyan Vangibhurathachhi

1809 Montgomery Ln , prosper Texas 75078.

Srinivasa.Kalyan2627@gmail.com

Abstract

Enterprise data management has passed through a series of architectural paradigms, each emerging to address the limitations of the one before it. This paper presents a structured comparative analysis of the three most influential approaches—the data warehouse, the data lake, and the data mesh—together with the transitional lakehouse pattern that bridges the first two. The study traces how enterprise architecture has shifted over three decades from centralised, highly structured repositories toward flexible and increasingly decentralised designs, and it compares the approaches systematically across seven dimensions: data structure, ownership model, scalability, governance, cost, suitability for business intelligence, and suitability for artificial intelligence and machine learning. Using a qualitative comparative methodology grounded in a review of the academic literature and documented industry practice, the paper finds that no single architecture is universally superior. The data warehouse offers consistency, reliability, and strong governance but handles unstructured data poorly and can become an organisational bottleneck. The data lake provides flexibility and low-cost storage for diverse data but frequently suffers from weak governance, giving rise to unreliable "data swamps." The data mesh addresses the organisational bottleneck of centralised platforms by decentralising data ownership to business domains while maintaining shared standards through federated governance, but it introduces significant coordination, skill, and governance challenges of its own. The analysis concludes that these paradigms are complementary as often as they are competing, and that most large organisations now operate hybrid architectures combining the reliability of the warehouse, the flexibility of the lake, and the organisational principles of the mesh. The appropriate choice depends on an organisation's size, sector, and data maturity rather than on adopting the latest trend. The broader trajectory reflects a shift in how organisations conceive of data—from a centralised asset to be stored toward a distributed product to be owned and shared.

Keywords: data warehouse; data lake; data mesh; lakehouse; enterprise architecture; data governance; comparative analysis; decentralisation

1. Introduction

1.1 Background context

The way organisations store, manage, and use data has changed profoundly over the past thirty years. As data has grown in volume, variety, and strategic importance, the architectures used to manage it have been repeatedly redesigned, and each redesign has responded to the shortcomings that became apparent in its predecessor. Understanding this progression is valuable not merely as a matter of historical interest but because many organisations today operate elements of all these architectures simultaneously and must decide how to combine or transition between them.

The data warehouse, which rose to prominence in the 1990s, established the idea of a single, consistent, centralised repository for structured data to support reporting and analysis. As organisations began to collect large volumes of unstructured and semi-structured data—text, images, logs, and sensor readings—the data lake emerged in the 2010s to provide cheap, flexible storage for raw data of any type. More recently, as even well-engineered central platforms began to strain under the demands of large, fast-moving organisations, the data mesh has proposed a fundamentally different solution: decentralising data ownership to the business domains that produce and understand the data, while maintaining shared governance and interoperability.

1.2 Problem statement

These approaches are frequently discussed in isolation, or presented as successive trends in which each replaces the last, and this framing can obscure both their relationships and their appropriate uses. Practitioners are often encouraged to adopt the newest paradigm without a clear understanding of the trade-offs involved or of whether their organisation is ready for it. The data mesh in particular is sometimes presented as a technology to be purchased rather than an organisational model to be adopted, leading to costly and unsuccessful implementations. There is a need for a balanced, structured comparison that clarifies what each approach offers, where each is appropriate, and how they relate to one another.

1.3 Research gap

While there is extensive literature on each architecture individually, and a growing number of practitioner accounts of data-mesh adoption, there are comparatively few systematic, dimension-by-dimension comparisons that place the warehouse, lake, and mesh side by side

using a consistent analytical framework. Much of the available material is either promotional, favouring a particular approach, or narrowly technical. This paper addresses that gap by providing a structured and balanced comparative analysis grounded in both scholarship and documented practice.

1.4 Research questions and significance

The study is guided by three questions: first, what drove the evolution from the data warehouse to the data lake to the data mesh; second, how do these architectures compare across a consistent set of technical and organisational dimensions; and third, how should organisations choose between them or combine them. Answering these questions matters because the choice of data architecture has far-reaching consequences for an organisation's agility, cost, governance, and ability to support analytics and AI. A poor choice—whether adopting a paradigm prematurely or clinging to one past its usefulness—can waste substantial investment and impede the organisation's use of its data. The remainder of the paper sets out the objectives and scope, reviews the literature, describes the comparative methodology, presents the comparative analysis with embedded figure descriptions, discusses the implications and guidance for selection, and concludes.

2. Objectives

The primary objective of this study is to compare the data warehouse, data lake, and data mesh architectures systematically across a consistent set of dimensions. The supporting objectives are: to explain the drivers behind the evolution from one paradigm to the next; to identify the strengths and limitations of each approach; to clarify how the transitional lakehouse pattern relates to the warehouse and lake; and to offer practical guidance on selecting among these architectures or combining them according to organisational context.

3. Scope of Study

This study examines enterprise data architectures at the level of their design principles, organisational implications, and trade-offs. It covers the data warehouse, data lake, lakehouse, and data mesh, and compares them across technical and organisational dimensions. It does not provide detailed implementation instructions, evaluate specific vendor products, or address the internal mechanics of database engines, which lie outside its scope. The analysis is comparative and conceptual rather than empirical; it synthesises established scholarship and documented

practice to compare the approaches, but it does not measure their performance through controlled experiments. The guidance offered is general and intended to be adapted to particular organisational contexts, recognising that the right choice depends on an organisation's size, sector, regulatory environment, and data maturity.

4. Literature Review

4.1 The data warehouse

The data warehouse concept was formalised through the work of Inmon (2005), who advocated an integrated, subject-oriented, enterprise-wide repository, and Kimball and Ross (2013), who developed dimensional modelling techniques optimised for querying and reporting. The warehouse brought discipline and consistency to enterprise reporting by integrating structured data from multiple operational systems into a single source of truth. Its strengths lay in reliability, data quality, and strong central governance, which made it well suited to regulated reporting and business intelligence. However, warehouses were costly to scale, required extensive upfront modelling, and were poorly equipped to handle the unstructured and semi-structured data that organisations increasingly accumulated.

4.2 The data lake and the lakehouse

The data lake emerged as a response to these limitations, offering low-cost storage for raw data of any format and supporting the large-scale analytics and machine learning associated with the big-data era (Khine & Wang, 2018). Lakes gave data scientists the flexibility to store and explore diverse data without imposing a schema in advance. This flexibility, however, came at a cost: without strong governance, metadata, and quality controls, many lakes degenerated into "data swamps" in which data was poorly documented, inconsistent, and difficult to trust (Nargesian et al., 2019). The difficulty of governing lakes and the desire to combine their flexibility with the reliability of warehouses led to the development of the lakehouse, a hybrid architecture that adds transactional guarantees, schema management, and reliable tables on top of low-cost lake storage (Armbrust et al., 2021). The lakehouse has become a widely adopted foundation because it supports both raw exploration and governed, reliable analytics within a single platform.

4.3 The data mesh

The data mesh, articulated by Dehghani (2022), represents a shift in focus from technology to organisation. Its central argument is that the principal bottleneck in large enterprises is not the

10.48047/jocaaa.2024.33.05.91

storage technology but the centralised data team responsible for serving the entire organisation, which cannot keep pace with the diverse and changing needs of many business domains. Drawing on principles from domain-driven design and product management, the mesh proposes four principles: domain-oriented ownership of data, treating data as a product, a self-serve data platform, and federated computational governance. Early adopters have reported improvements in scalability and agility, but researchers have also documented significant challenges, including the need for new skills within domains, cultural resistance, unclear incentives, and the difficulty of maintaining consistent standards across autonomous teams (Machado et al., 2022; Bode et al., 2023). The literature broadly agrees that the data mesh is an organisational transformation as much as a technical one, and that its success depends heavily on organisational maturity.

4.4 Synthesis of the literature

Taken together, the literature describes an evolution from centralised structure toward flexible and then decentralised designs, with hybrid models becoming increasingly common. Each paradigm arose to solve the problems of its predecessor while introducing new challenges of its own. What the literature lacks, and what this paper provides, is a systematic side-by-side comparison that allows the trade-offs to be weighed consistently and that supports informed choice rather than trend-following.

5. Research Methodology

5.1 Research approach

This study employs a qualitative comparative methodology, which is well suited to clarifying the principles and trade-offs of competing architectural approaches. Because the subject is architectural and organisational rather than narrowly empirical, the aim is to compare the approaches systematically and transparently rather than to test them statistically. Comparative conceptual analysis of this kind is a recognised method in information-systems research, where the goal is to produce structured understanding and practical guidance.

5.2 Data sources and selection

The analysis draws on two bodies of evidence. The first is a structured review of the academic literature, including foundational texts and peer-reviewed articles published between 2005 and 2023, identified through academic databases and selected for direct relevance to the architectures under study. The second is an examination of documented industry practice,

including widely adopted reference architectures and published accounts of implementation experience. Combining scholarly and practitioner sources ensures that the comparison reflects both established theory and real-world application.

5.3 Comparison framework

A consistent set of seven comparison dimensions was defined in advance to structure the analysis: data structure (the types of data the approach handles well), ownership model (centralised or decentralised responsibility), scalability (technical and organisational), governance (the strength and mode of control), cost (storage and operational), suitability for business intelligence, and suitability for artificial intelligence and machine learning. These dimensions were chosen because they capture the principal technical and organisational differences that matter to organisations making architectural decisions. Each architecture was then assessed against every dimension, drawing on the reviewed sources so that the judgements rest on established understanding rather than opinion.

5.4 Limitations of the methodology

Because the methodology is comparative and conceptual rather than experimental, its conclusions are analytical rather than measured. The assessments on each dimension reflect the weight of documented experience and scholarship, but they are necessarily general and may vary in specific implementations. Empirical comparison—measuring cost, agility, and governance outcomes across organisations using different architectures—would complement this analysis and is noted as a direction for future research.

6. Comparative Analysis

6.1 The evolution as an accumulation of principles

Before comparing the architectures dimension by dimension, it is important to understand their relationship over time, because they are often misunderstood as a simple sequence of replacements.




Figure 1. Evolution timeline of enterprise data architectures.

6.2 Data structure

The architectures differ markedly in the types of data they handle well. The data warehouse is optimised for structured data organised around predefined schemas and handles unstructured data poorly. The data lake, by contrast, accepts data of any format, from structured tables to raw text and images, making it far more flexible but also harder to govern. The data mesh is format-agnostic, because it is primarily an organisational model layered over various underlying storage technologies, including warehouses and lakehouses; its concern is ownership and interoperability rather than storage format. The lakehouse occupies a middle position, supporting both raw and structured data within a single platform.

6.3 Ownership model

Ownership is the dimension on which the data mesh differs most sharply from its predecessors. In the warehouse and the lake, data is typically owned and managed centrally by a single data team responsible for the entire organisation. This centralisation provides control and consistency but concentrates responsibility in one team. The data mesh, by contrast, distributes ownership to the business domains that produce and use the data, so that each domain is responsible for its own data end to end. This decentralisation is the defining characteristic of the mesh and the source of both its advantages and its challenges.




Figure 2. Centralised versus decentralised ownership topology.

6.4 Scalability

Scalability differs in kind as well as degree across the architectures. Warehouses can scale technically to large volumes but tend to become organisational bottlenecks, because all changes and new requirements flow through a single central team and modelling process. Lakes scale cheaply in terms of storage but can become unmanageable without governance, so that their technical scalability is undermined by a loss of trust and usability. The data mesh is explicitly designed to scale organisationally, allowing many domain teams to work in parallel without depending on a central team; however, this organisational scalability depends on strong shared infrastructure and governance, without which the mesh fragments. The decisive point is that the mesh addresses a different kind of scalability—organisational rather than purely technical—than its predecessors.

6.5 Governance

Governance also varies significantly. The warehouse offers strong central governance, with consistent standards, quality controls, and access management applied by the central team, which is one of its principal strengths and the reason it remains favoured in regulated contexts. The lake often lacks governance unless it is deliberately added, which is the root cause of the data-swamp problem. The data mesh uses federated governance, in which global standards are defined centrally by a cross-domain group but applied and owned locally by each domain, ideally embedded in the shared platform as automated rules. This approach is powerful in principle because it balances consistency with autonomy, but it is demanding in practice, requiring both platform investment and organisational discipline.

6.6 Cost, business intelligence, and AI suitability

The remaining dimensions can be considered together. In terms of cost, warehouses are relatively expensive owing to their proprietary engines and modelling effort; lakes are inexpensive to store data in but can become costly to govern and make usable; and meshes are efficient at scale but require significant upfront investment in platform and culture. For business intelligence, the warehouse remains excellent, providing fast, reliable, structured reporting; the lake is less suited to standard reporting without additional structure; and the mesh depends on the quality of the underlying domain data products. For artificial intelligence and machine learning, the lake and the mesh are generally better suited than the warehouse, because they handle diverse data and support the flexible access that machine learning requires; the warehouse's rigid structure is a limitation for AI workloads.

Dimension	 Data Warehouse (Centralised, structured)	 Data Lake (Centralised, multi-structured)	 Data Mesh (Decentralised, domain-driven)
 Data structure supported	 Moderate Strong for structured data (tables, relational).	 Strong Supports structured, semi-structured and unstructured data (files, logs, images, etc.).	 Strong Supports multi-structured data across domains.
 Ownership model	 Moderate Central IT owns and manages.	 Moderate Centralised ownership (often IT).	 Strong Domain-oriented ownership (across business units).
 Scalability	 Moderate Scales vertically; higher cost at large scale.	 Strong Scales horizontally at low cost (cloud storage).	 Strong Organisational and technical scalability across domains.
 Governance	 Strong Mature governance, data quality and compliance.	 Weak Governance is more challenging due to diverse and unstructured data.	 Moderate Requires federated governance; higher effort and organisational maturity.
 Cost	 Weak Higher storage and compute cost.	 Strong Low-cost storage (especially cloud).	 Moderate Cost varies; requires investment in federated platforms and governance.
 Suitability for Business Intelligence	 Strong Highly suitable (optimised for BI and reporting).	 Moderate Suitable, but requires additional data processing and modelling.	 Moderate Suitable with domain-specific data products; may require integration.
 Suitability for AI and Machine Learning	 Weak Limited due to structured data and schema constraints.	 Strong Highly suitable (access to large, raw and diverse data).	 Strong Highly suitable with domain-specific AI/ML data products.
 Overall summary	Best for governed, structured data and traditional BI.	Best for flexible, large-scale, low-cost storage and AI/ML.	Best for organisational scalability, domain ownership and advanced analytics (AI/ML), but requires higher governance effort.
Legend (strength of fit)  Strong (well suited)  Moderate (partially suited)  Weak (least suited) Key takeaways • Data warehouse: strong governance and BI, but limited for unstructured data and AI. • Data lake: flexible and cost-effective, but governance is challenging. • Data mesh: scalable and AI-ready, but requires higher governance effort and organisational maturity.			

Figure 3. Architectural comparison across key dimensions.

6.7 A key finding: architectures as complementary rather than successive

A central finding of the comparison is that these architectures are not strictly sequential replacements. Many organisations run warehouses and lakes together, and the data mesh is frequently implemented on top of existing lakehouse infrastructure rather than replacing it. The evolution is therefore better understood as an accumulation of principles than as a series of clean substitutions. Each paradigm contributes a capability that remains useful: the warehouse

contributes reliability and structure, the lake contributes flexibility and low-cost storage, and the mesh contributes distributed ownership and organisational scalability. This finding has direct implications for how organisations should approach architectural decisions, which are discussed below.

7. Discussion

7.1 Interpretation of the comparison

The comparative analysis shows that the shift from warehouse to lake to mesh reflects two underlying trends. The first is a movement from structured to flexible data handling, driven by the growing diversity of data and the demands of analytics and AI. The second is a movement from centralised to decentralised ownership, driven by the organisational reality that a single central team cannot serve the data needs of a large, fast-moving enterprise. These two trends are distinct: the first is primarily technical, while the second is primarily organisational, and the data mesh is best understood as a response to the second rather than the first.

7.2 Guidance on selection

Importantly, the analysis suggests that no architecture is universally best, and that the appropriate choice depends on organisational context. The warehouse remains the right choice where structured reporting and strong central control are paramount, such as in finance, compliance, and regulated reporting. The lake, or more commonly the lakehouse, suits organisations that need flexibility and work heavily with unstructured data or machine learning. The data mesh suits large, complex organisations in which central teams have become bottlenecks and in which domains are mature enough to own their data responsibly. For smaller organisations, or those with limited data maturity, the mesh is likely to add unnecessary complexity and may fragment standards rather than improve them.

7.3 The prevalence of hybrid architectures

This contextual guidance explains why hybrid architectures dominate in practice. A large enterprise might maintain a warehouse for regulated reporting, a lakehouse for analytics and AI, and mesh principles to distribute ownership across domains, all at the same time. The decisive factor in whether such arrangements succeed, particularly where the mesh is involved, is organisational maturity rather than technology. The mesh depends less on any specific tool than on culture, skills, incentives, and governance discipline. Where these are lacking, decentralisation can fragment standards and reduce trust in data, producing worse outcomes

than a well-run central platform. Organisations should therefore assess their readiness honestly before adopting decentralised models.

7.4 Limitations and future research

The principal limitation of this study is that its conclusions rest on documented experience and scholarship rather than on controlled measurement. The assessments on each dimension are general and may differ in particular implementations. Future research should complement this conceptual comparison with empirical studies measuring cost, agility, governance outcomes, and AI-deployment success across organisations using different architectures, and with longitudinal case studies examining how organisations transition between paradigms and combine them over time.

8. Conclusion

The evolution from data warehouses to data lakes to data meshes represents a steady shift from centralised, structured repositories toward flexible and decentralised designs. Each paradigm arose to address the limitations of its predecessor, and each contributes principles that remain valuable: the warehouse's reliability and governance, the lake's flexibility and low-cost storage, and the mesh's distributed ownership and organisational scalability. The comparative analysis presented here, conducted across seven consistent dimensions, shows that these approaches are complementary as often as they are competing, and that most large organisations now combine them in hybrid architectures. No single approach is universally superior; the appropriate choice depends on an organisation's size, sector, regulatory environment, and data maturity rather than on following the latest trend. Above all, the mesh should be understood as an organisational transformation requiring genuine readiness, not a technology to be purchased. The broader trajectory of enterprise data architecture reflects a fundamental change in how organisations conceive of data—moving from a centralised asset to be stored toward a distributed product to be owned, shared, and trusted across the enterprise.

References

- 1: **A. Mahendra Vardhan and S. Sridhar**, "Determining False Positive Analysis of Software Vulnerabilities with Predefined Scan Rules using Random Forest Classifier and Decision Tree Technique," 2022 4th International Conference on Advances in Computing, Communication Control and Networking (ICAC3N), Greater Noida, India, 2022, pp. 622-625,
DOI: <https://doi.org/10.1109/ICAC3N56670.2022.10074458>

10.48047/jocaaa.2024.33.05.91

2: Dip Bharatbhai Pate, Dynamic vs. Static visualizations: Understanding their use cases in big data analysis, International Journal of Science and Research Archive, 2022, 05(02), 392-395., DOI <https://doi.org/10.30574/ijrsra.2022.5.2.0077>

3: Dip Bharatbhai Patel, Building scalable business intelligence systems in the cloud: A technical approach, International Journal of Science and Research Archive, 2021, 02(01), 212-215., Received on 03 February 2021; revised on 17 April 2021; accepted on 19 April 2021, DOI: <https://doi.org/10.30574/ijrsra.2021.2.1.0047>

4: Dip Bharatbhai Patel, Understanding the Impact of Text Analytics on Social Media Sentiment Analysis, Vol. 5 No. 5 (2019) ,September-October, DOI: <https://doi.org/10.32628/CSEIT1936413>

5: Dip Bharatbhai Patel, Real-Time Data Streaming in BI Dashboards: Opportunities and Challenges, Archives / Vol. 5 No. 4 (2019) : July-August, DOI: <https://doi.org/10.32628/CSEIT192142>

6: Dip Bharatbhai Patel, the Role of Data Visualization in Modern Intelligence, **Archives / Vol. 2 No. 2 / Research Article, Published, 2020-11-28 DOI:** <https://doi.org/10.32996/jcsts.2020.2.2.4>

7: A. Srivastava, "AI-assisted cloud migration of healthcare claims data from legacy systems: A metadata-driven framework," Int. J. Commun. Networks Inf. Secur., vol. 14, no. 3, Nov. 2022. [Online]. Available: <https://ijcnis.org/index.php/ijcnis/article/view/8682>

8: A. Srivastava, "Enhancing provider and claims data accuracy using artificial intelligence on cloud-native data platforms," J. Comput. Anal. Appl., vol. 31, no. 4, Nov. 2023. [Online]. Available: <https://eudoxuspress.com/index.php/pub/article/view/4347/3189>

9: A. Srivastava, "Optimizing large-scale data migration for cloud-native architectures," J. Comput. Anal. Appl., vol. 31, no. 2, pp. 652–662, May 2023. [Online]. Available: <https://eudoxuspress.com/index.php/pub/article/view/4603>

10: Backup Integrity and Recovery Readiness Assessment for High-Availability Databases
Raghu Gollapudi, Computer Fraud & Security, Vol. 2023, Issue 8, pp. 35-48, [10.52710/cfs.1031](https://doi.org/10.52710/cfs.1031)

11: Risk-Controlled Near-Zero-Downtime Oracle Database Migration Using GoldenGate
Raghu Gollapudi, International Journal of Computational and Experimental Science and Engineering, 8(3), 2022 ,[10.22399/ijcesen.5382](https://doi.org/10.22399/ijcesen.5382)

12: Manoj Kumar Kagitha, Automated Security Compliance in Ci/Cd: A Natural Language Processing (NLP) Approach to Detecting Vulnerabilities in Infrastructure-As-Code (Terraform/Bicep), Publication year 2021, Volume 2021, Issue 12, <https://computerfraudsecurity.com/index.php/journal/article/view/970>

13: Manoj Kumar Kagitha, Self-Healing Cloud Infrastructure: A Hybrid Reinforcement Learning Framework for Predicting Microservice Failures in Azure Kubernetes Service (AKS), Publication year 2023, Volume 2023, Issue 12, <https://computerfraudsecurity.com/index.php/journal/article/view/971>

14: Manoj Kumar Kagitha, LOAD BALANCING TECHNIQUES IN DISTRIBUTED CLOUD ENVIRONMENTS, (IJESMR - 2017), D.O.I:10.53555/jcr.v6:i1.13380, <https://jcreview.com/archives/volume-6/issue-1/13380>

15: Manoj Kumar Kagitha, AI-DRIVEN OBSERVABILITY FOR HYPERSCALE COLOCATION DATA CENTERS: A CASE STUDY OF LOG ANALYTICS AND ANOMALY DETECTION

PIPELINES, Archives / Vol. 27 No. 7 (2019):

JOCAAA, <https://www.eudoxuspress.com/index.php/pub/article/view/5053>

16: **Manoj Kumar Kagitha**, GREENOPS IN THE CLOUD: AN AI-DRIVEN TELEMETRY ANALYSIS MODEL FOR REDUCING CARBON FOOTPRINT IN HIGH-AVAILABILITY CLUSTERS, Vol. 28 No. 6 (2020): JOCAAA,

<https://www.eudoxuspress.com/index.php/pub/article/view/5054>

17: **Naresh Lokiny. (2022)**. Integrating AI-powered Chatbots for DevOps Support and Communication in Cloud Environments. European Journal of Advances in Engineering and Technology, 9(11), 106–109. <https://doi.org/10.5281/zenodo.13325989>

18: **Naresh Lokiny, (2021)**, "Disaster Recovery and Business Continuity Planning in DevOps Cloud with AI", International Journal of Science and Research (IJSR), 10(3), 2023-2027. <https://dx.doi.org/10.21275/SR24724151733>,

<https://www.ijsr.net/getabstract.php?paperid=SR24724151733>

19: **Naresh Lokiny, & Ranganath Nandanampati.** (2020). DevSecOps: Integrating Security into DevOps with AI in Cloud. Journal of Scientific and Engineering Research, 7(10), 239–242. <https://doi.org/10.5281/zenodo.13348695>

20: **Naresh Lokiny, & Pradip Reddy.** (2021). Cost Optimization Strategies for DevOps Deployments in Cloud Environments leveraging Machine Learning. European Journal of Advances in Engineering and Technology, 8(3), 69–72. <https://doi.org/10.5281/zenodo.13325845>

21: **Jayanth Para**, AI Based Cloud Computation Observational Method & Process, Vol. 51 No. 4 (2023): October-December 2023, Power System Protection and Control, ISSN-1674-3415,

<https://pspac.info/index.php/dlbh/article/view/171> , DOI: <https://doi.org/10.46121/pspc.51.4.3>

22: **Jobayar Alom, Ahsan Ahmed.** (2023). Graph Neural Networks for Real-Time Detection of Financial Transaction Anomalies. Acta Scientiae, 24(5), 82–90. <https://doi.org/10.22178/acta.24.5.6>

23: **Kaleshwar Aryasomayajula**, PREVENTING BIAS IN AI-BASED DECISION-MAKING: ANALYZING TECHNIQUES TO REMOVE UNFAIR PREJUDICE IN ALGORITHMIC OUTCOMES, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/294>

24: **Kaleshwar Aryasomayajula**, MODERN DEVELOPMENT PRACTICES: INVESTIGATIONS INTO SERVERLESS COMPUTING, LOW-CODE/NO-CODE PLATFORMS, AND DEVELOPER PRODUCTIVITY IN REMOTE ENVIRONMENTS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/364>

25: **Vikas Suresh Bagora**, KERNEL-LEVEL PERFORMANCE OPTIMIZATION FOR CARRIER-GRADE BROADBAND AND HIGH-SPEED NETWORKING SYSTEMS, Vol. 47 No. 1 (2019), Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/359>

26: **Vikas Suresh Bagora**, SCALABLE 128G FIBRE CHANNEL AND ETHERNET CONVERGENCE FOR NEXT-GENERATION DATA CENTER NETWORKS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/362>

27. **Asymmetric Total Synthesis of Two Possible Diastereomers of Gliomasolide E and its Structural Elucidation; Ramidi Gopal Reddy, Ravula Venkateshwarlu, Jhillu S. Yadav and Debendra K. Mohapatra.** (J. Org. Chem. 2017, 82(2), 1053-1063). <https://doi.org/10.1021/acs.joc.6b02611> , <https://pubs.acs.org/doi/abs/10.1021/acs.joc.6b02611>

10.48047/jocaaa.2024.33.05.91

- 28: **Amanullakhan Pathan Jayadip Ghanshyambhai Tejani, Rahul Shah, Hiral Vaghela, Shailesh, Vajapara , Controlled Synthesis and Characterization of Lanthanum Nanorods**, 2020/5/1, International Journal of Thin Films Science and Technology, Natural Sciences Publishing Corporation (NSP)
https://scholar.google.com/citations?view_op=view_citation&hl=en&user=efbNMkwAAAAJ&citation_for_view=efbNMkwAAAAJ:M3NEmzRMiKIC
- 29: **Stereoselective synthesis of the C3-C15 fragment of callispongolide; Ramidi Gopal Reddy, Jhillu S. Yadav and Debendra K. Mohapatra. (Tetrahedron Letters. 2018, 59, 3579-3582).**
<https://doi.org/10.1016/j.tetlet.2018.08.041>,
<https://www.sciencedirect.com/science/article/pii/S0040403918310426>
30. **Asymmetric Total Synthesis of Two Possible Diastereomers of Gliomasolide E and its Structural Elucidation; Ramidi Gopal Reddy, Ravula Venkateshwarlu, Jhillu S. Yadav and Debendra K. Mohapatra. (J. Org. Chem. 2017, 82(2), 1053-1063).**
<https://doi.org/10.1021/acs.joc.6b02611>, <https://pubs.acs.org/doi/abs/10.1021/acs.joc.6b02611>
- 31: **Kaleshwar Aryasomayajula, PREVENTING BIAS IN AI-BASED DECISION-MAKING: ANALYZING TECHNIQUES TO REMOVE UNFAIR PREJUDICE IN ALGORITHMIC OUTCOMES**, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/294>
- 32: **Kaleshwar Aryasomayajula, MODERN DEVELOPMENT PRACTICES: INVESTIGATIONS INTO SERVERLESS COMPUTING, LOW-CODE/NO-CODE PLATFORMS, AND DEVELOPER PRODUCTIVITY IN REMOTE ENVIRONMENTS**, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/364>
- 33: **Kaleshwar Aryasomayajula, Micro services: “A Comparative Analysis of Monolithic vs. Microservice Architectures in High-Scalability Cloud Environments.**, Vol. 51 No. 2 (2023): **April-June 2023** , Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/374>
- 34: **Controlled Synthesis and Characterization of Lanthanum Nanorods, Author(s), Jayadeep Tejani, Rahul Shah, Hiral Vaghela, Shailesh Vajapara, Amanullakhan Pathan,** doi:10.18576/ijtfst/090205, URL: <https://www.naturalspublishing.com/Article.asp?ArtcID=20705>, May-2020
- 35: **Sudipkumar Ghanvat , Aishwarya Badlani, Shreya Sandeep Joshi, Marketing Effectiveness in the Post-Cookie Era: A Comparative Analysis of First Party and Zero-Party Data Strategies on Campaign Performance and Customer Equity**, Vol. 31 No. 4 (2023): JOCAAA , <https://www.eudoxuspress.com/index.php/pub/article/view/3940>
- 36: **Chander Vijay S Sanbhi, BAYESIAN UPDATING OF RAM MODELS FOR DYNAMIC AVAILABILITY FORECASTING UNDER REALISTIC DOWNTIME CONSTRAINTS**, Vol. 51 No. 3 (2023): July-September 2023, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/381>
- 37: **Chander Vijay S Sanbhi, HYBRID RELIABILITY MODELLING FOR ROTATING EQUIPMENT: PHYSICS-INFORMED LEARNING WITH SPARSE FAILURE DATA**, Vol. 51 No. 4 (2023): October-December 2023, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/382>

10.48047/jocaaa.2024.33.05.91

38: **A. Srivastava**, "AI-assisted cloud migration of healthcare claims data from legacy systems: A metadata-driven framework," *Int. J. Commun. Networks Inf. Secur.*, vol. 14, no. 3, Nov. 2022. [Online]. Available: <https://ijcnis.org/index.php/ijcnis/article/view/8682>

39: **A. Srivastava**, "Enhancing provider and claims data accuracy using artificial intelligence on cloud-native data platforms," *J. Comput. Anal. Appl.*, vol. 31, no. 4, Nov. 2023. [Online]. Available: <https://eudoxuspress.com/index.php/pub/article/view/4347/3189>

40: **A. Srivastava**, "Optimizing large-scale data migration for cloud-native architectures," *J. Comput. Anal. Appl.*, vol. 31, no. 2, pp. 652–662, May 2023. [Online]. Available: <https://eudoxuspress.com/index.php/pub/article/view/4603>

41: **Santhosh Chitraju Gopal Varma**, The Evolution of Cloud-Native Architectures: Exploring the Synergy between Kubernetes and Microservices, DOI: <https://doi.org/10.63282/3050-9246.IJETCSIT-V1I4P104>, Vol. 1 No. 4 (2020), <https://ijetcsit.org/index.php/ijetcsit/article/view/108>

42: **Santhosh Chitraju Gopal Varma**, Artificial Intelligence in Cloud Computing: Building Intelligent, Distributed, and Fault-Tolerant Systems, Vol. 3 No. 1 (2022), DOI: <https://doi.org/10.63282/3050-9416.IJAIBDCMS-V3I1P105>, <https://ijaibdcms.org/index.php/ijaibdcms/article/view/78>

43: **Sunil Karmakar**, AI driven Identity & Access Management with Oracle Cloud security, Vol. 30 No. 2 (2022): JOCAAA, <https://eudoxuspress.com/index.php/pub/article/view/4118>