

Designing Enterprise Data Architecture for AI-First Institutions

Srinivasa Kalyan Vangibhurathachhi

1809 Montgomery Ln , prosper Texas 75078.
Srinivasa.Kalyan2627@gmail.com

Abstract

Artificial intelligence has moved from isolated experiments to the operational core of many organisations, giving rise to a new category of institution in which AI drives decisions, products, and services rather than supplementing them. For these "AI-first" institutions, the enterprise data architecture is no longer a back-office concern but a strategic foundation that determines how quickly and reliably intelligent systems can be built and deployed. Yet most enterprise architectures in use today were designed for reporting and transaction processing, and they struggle to meet the volume, variety, velocity, and reproducibility requirements of machine learning. This paper examines the principles, components, and organisational conditions required to design enterprise data architecture for AI-first institutions. Using a qualitative, design-oriented methodology that combines a structured literature review with an analysis of established industry reference architectures, the study develops and justifies a five-layer reference model comprising ingestion, storage, feature processing, model serving, and a cross-cutting governance layer. It argues that AI-first architecture differs from conventional architecture not only in scale but in orientation: data must be treated as a continuously flowing, reusable product rather than a static historical record, and governance must be embedded throughout rather than applied at the end. The analysis finds that the most decisive factors in successful AI-first architecture are not raw technology but reproducibility, timeliness, data quality, and governance discipline, supported by an organisational culture of data ownership. Institutions that align these elements shorten the path from raw data to deployed model and reduce the risk of unreliable automated decisions, while those that neglect governance build systems that are fast but fragile. The paper offers a practical reference model and a set of design principles that institutions can adapt to their own scale, sector, and maturity.

Keywords: enterprise data architecture; artificial intelligence; AI-first institutions; data pipelines; feature store; MLOps; data governance; data-as-a-product

1. Introduction

1.1 Background context

Artificial intelligence has become a central driver of competitiveness across almost every sector. Banks use machine learning for credit scoring and fraud detection, hospitals for diagnostic support and patient-risk prediction, retailers for demand forecasting and personalisation, and public institutions for resource allocation and service delivery. This widespread adoption has produced a distinctive kind of organisation, often described as "AI-first," in which intelligent systems are embedded in core processes and products rather than treated as peripheral tools. In such institutions, the quality, timeliness, and availability of data directly determine the quality of the AI, and therefore the quality of the organisation's decisions and offerings. The data architecture that supplies these systems has consequently moved from a technical support function to a matter of strategic importance.

The term data architecture refers to the overall structure of an organisation's data assets and the systems, processes, and policies that govern how data is collected, stored, transformed, served, and protected. For decades, enterprise data architecture was shaped by the needs of business intelligence: producing accurate reports, supporting compliance, and enabling structured analysis of historical events. The dominant designs—relational databases and data warehouses—were optimised for these purposes, with data cleaned and loaded in regular batches and organised around predefined schemas. These systems served their intended purpose well, but they were not designed for the demands that machine learning now places on data.

1.2 Problem statement

The central problem addressed by this paper is that conventional enterprise data architectures are poorly suited to the requirements of AI-first institutions, and this mismatch has become a significant barrier to scaling AI. Machine learning requires large volumes of diverse data, including unstructured data such as text, images, and sensor readings; it requires rapid and often real-time access to data; and, critically, it requires the ability to reproduce the exact data used to train a model so that results can be audited and models retrained reliably. Traditional architectures, built around batch loading of structured data into fixed schemas, cannot easily meet these needs. As a result, many organisations find that although they can build promising

AI prototypes, they cannot move them into reliable production use because the underlying data foundation does not support it. This gap between experimentation and production is widely recognised as one of the main reasons AI initiatives fail to deliver value.

1.3 Research gap

While there is extensive literature on individual components of AI infrastructure—such as feature stores, data pipelines, and machine-learning operations (MLOps)—and a growing body of work on data architecture paradigms such as data lakes and data meshes, there is relatively little consolidated guidance that brings these elements together into a coherent architectural model specifically designed for AI-first institutions. Much of the existing material is either highly technical and focused on single tools, or conceptual and focused on organisational principles, with few sources integrating the technical and organisational dimensions. This paper addresses that gap by synthesising both into a single reference model and set of design principles.

1.4 Research questions and significance

The study is guided by three questions: first, how do the data requirements of AI-first institutions differ from those of traditional reporting-oriented organisations; second, what components and layers constitute an enterprise data architecture capable of meeting those requirements; and third, what organisational and governance conditions are necessary for such an architecture to succeed. Answering these questions matters because the ability to deploy AI reliably and at scale increasingly determines organisational performance, and because poorly designed data foundations lead not only to wasted investment but to automated decisions built on flawed or unreliable data. The remainder of the paper sets out the study's objectives and scope, reviews the relevant literature, describes the methodology, presents the proposed reference architecture with embedded figure descriptions, discusses the implications and trade-offs, and concludes with practical guidance and directions for future work.

2. Objectives

The primary objective of this study is to define the design principles and components of an enterprise data architecture suited to AI-first institutions. The supporting objectives are: to compare the data requirements of AI-first institutions with those of traditional data architectures; to propose a layered reference model that organisations can adapt to their own context; to clarify the role of governance and data-as-a-product thinking in supporting AI at

scale; and to identify the organisational and cultural conditions that influence whether such an architecture succeeds in practice.

3. Scope of Study

This study focuses on the architectural and organisational design of enterprise data systems that support artificial intelligence and machine learning. It covers the main architectural layers—ingestion, storage, feature processing, model serving, and governance—and the organisational conditions that enable them. It does not address the detailed design of specific machine-learning algorithms, the selection of particular vendor products, or the mathematical foundations of model training, which lie outside its concern. The study is conceptual and design-oriented rather than empirical; it synthesises existing knowledge and established practice to construct and justify a reference model, but it does not test that model through implementation in a specific institution. The guidance offered is intended to be adaptable across sectors and organisation sizes, though the paper notes throughout that the appropriate design depends on an institution's scale, sector, regulatory environment, and data maturity.

4. Literature Review

4.1 The evolution of enterprise data architecture

Enterprise data architecture has evolved through several distinct stages, each responding to the limitations of its predecessor. The relational database, introduced in the 1970s, established the foundation for structured data management. The data warehouse, formalised by Inmon (2005) and refined by Kimball and Ross (2013), brought the ability to integrate structured data from multiple sources into a single consistent repository optimised for querying and reporting. Warehouses became the backbone of enterprise business intelligence, but they were expensive to scale and ill-suited to the unstructured and semi-structured data that organisations increasingly collected.

The data lake emerged in the 2010s as a response to these limitations, offering low-cost storage for raw data of any format and enabling the large-scale analytics and machine learning associated with the big-data era (Khine & Wang, 2018). Lakes gave data scientists the flexibility to work with diverse data, but without strong governance many degenerated into "data swamps," in which data was poorly documented, inconsistent, and difficult to trust (Nargesian et al., 2019). This problem led to the development of the "lakehouse," a hybrid

design that combines the low-cost flexibility of the lake with the structure, reliability, and transactional guarantees of the warehouse (Armbrust et al., 2021). The lakehouse has become a common foundation for AI-oriented architectures because it supports both raw and curated data within a single platform.

4.2 The distinctive requirements of machine learning

The rise of machine learning has introduced requirements that earlier architectures did not anticipate. In an influential analysis, Sculley et al. (2015) described the "hidden technical debt" in machine-learning systems, observing that only a small fraction of the effort in deploying AI lies in the modelling itself; the majority lies in data collection, feature engineering, configuration, dependency management, and monitoring. This insight reframed the challenge of AI as primarily a data and systems-engineering problem rather than an algorithmic one, and it motivated the development of specialised infrastructure.

Two developments are particularly important. The first is the feature store, a centralised system for defining, computing, storing, and serving the engineered variables ("features") that models consume. Feature stores address a critical requirement of machine learning—reproducibility—by ensuring that the same feature definitions are applied consistently during both training and prediction, preventing the subtle discrepancies known as training-serving skew. The second is the set of practices known as MLOps, which apply the discipline of software engineering and operations to the machine-learning lifecycle, bringing version control, automated testing, continuous deployment, and monitoring to models and the data they depend on (Kreuzberger et al., 2023). Together, these developments reflect a recognition that AI requires not just storage and compute but a managed, reproducible lifecycle.

4.3 Organisational and governance perspectives

A further strand of literature emphasises that the challenges of AI-first architecture are as much organisational as technical. The data mesh concept, articulated by Dehghani (2022), argues that in large enterprises the bottleneck is not storage technology but the centralised data team responsible for serving the entire organisation, which cannot keep pace with the diverse needs of many domains. The mesh proposes decentralising data ownership to domain teams while maintaining shared standards through federated governance, and treating data as a product with clear ownership and quality guarantees. Other work stresses the importance of data quality and lineage, noting that AI amplifies the consequences of poor data because errors and biases propagate directly into automated decisions (Gudivada et al., 2017). Across this literature a

consistent theme emerges: supporting AI at scale depends on treating data as a governed, reusable product with clear ownership, rather than as an incidental by-product of operations. However, few sources integrate these organisational principles with the technical components into a single coherent architecture, which is the contribution this paper seeks to make.

5. Research Methodology

5.1 Research approach and design

This study uses a qualitative, design-oriented methodology appropriate to architectural research, where the objective is to produce a well-reasoned and justified design rather than to test a hypothesis through statistical analysis. Rather than collecting new empirical data, the study synthesises existing knowledge from scholarship and documented practice to construct a reference model. This "design-and-evaluate" approach is widely used in information-systems and software-architecture research, where proposed designs are developed from established principles and then assessed conceptually against known requirements.

5.2 Data sources and selection

The study draws on two main sources of evidence. The first is a structured review of the academic and technical literature on data architecture, machine-learning infrastructure, and data governance, covering peer-reviewed articles, foundational texts, and technical reports published between 2005 and 2023. Sources were identified through academic databases and reputable technical repositories and were selected for their direct relevance to AI workloads, data architecture, and governance. The second source is an analysis of the published reference architectures of major cloud and technology providers, which document widely adopted patterns and components. Examining these practitioner architectures alongside the scholarly literature ensures that the resulting model reflects both theoretical understanding and real-world practice.

5.3 Analytical procedure

The analysis proceeded in three stages. First, the literature and reference architectures were reviewed to identify the recurring components and requirements of AI-oriented data systems. Second, these components were grouped thematically into architectural layers, each corresponding to a distinct function in the flow of data from source to deployed model. Third, the resulting layered model was evaluated conceptually against the four defining requirements of AI workloads identified in the review—scale, timeliness, reproducibility, and governance—

to confirm that each layer addresses a genuine need and that the model as a whole is coherent and complete.

5.4 Limitations of the methodology

Because the study is conceptual rather than empirical, the proposed reference model is justified by reasoning and established practice rather than validated through implementation and measurement in a specific institution. This is an appropriate and common approach for architectural research, but it means that the conclusions are analytical rather than experimental. Empirical validation—measuring outcomes such as time-to-deployment, model reliability, and governance incidents in institutions that adopt the model—would be a valuable direction for future work and is noted as such in the discussion.

6. Findings and the Proposed Reference Architecture

6.1 Overview of the five-layer model

The analysis identifies five layers that together constitute an enterprise data architecture suited to AI-first institutions: ingestion, storage, feature processing, model serving, and a cross-cutting governance layer. The first four layers describe the vertical flow of data from raw sources to deployed models, while the fifth spans all of them to ensure quality, control, and accountability throughout.




Figure 1. Five-layer reference architecture for AI-first institutions.

6.2 The ingestion layer

The ingestion layer is responsible for collecting data from the institution's many sources and bringing it into the architecture. In contrast to traditional architectures, which typically rely on scheduled batch loading, the ingestion layer of an AI-first architecture must support both batch and streaming ingestion. Batch ingestion remains appropriate for large, periodic data transfers, but many AI applications—fraud detection, recommendation, dynamic pricing—depend on acting upon fresh data within seconds or minutes, which requires streaming ingestion.

Designing this layer to handle both modes is one of the first distinguishing features of an AI-ready architecture, and it directly addresses the requirement of timeliness.

6.3 The storage layer

The storage layer holds the institution's data in a form that balances cost, flexibility, and reliability. The analysis finds that the lakehouse pattern is well suited to AI-first institutions because it allows raw and unstructured data to be stored cheaply in a data lake while still providing structured, reliable, queryable tables for curated datasets. This combination is important because machine learning often requires access to raw data for exploration and feature engineering, as well as to clean, governed datasets for reliable training and auditing. The storage layer therefore addresses the requirement of scale, enabling the institution to retain large volumes of diverse data without prohibitive cost.

6.4 The processing and feature layer

The processing and feature layer is where raw data is transformed into the features that models consume, and it is the layer most distinctive to AI-first architecture. At its centre is the feature store, which allows features to be defined once and reused across many models and teams, and which ensures that the same transformations are applied during both training and prediction. This directly addresses the requirement of reproducibility, which is essential both for reliable model performance and for auditability. By centralising feature definitions, this layer also reduces duplication and inconsistency, which are common sources of error and inefficiency in large organisations.




Figure 2. From raw data to deployed model — the data lifecycle with governance checkpoints.

6.5 The model serving layer

The model serving layer supplies data to models running in production and manages their deployment, monitoring, and retraining. This layer embodies MLOps practices, treating models as managed assets with version control, automated testing, and continuous monitoring of performance. It ensures that models can be updated reliably as data changes and that degradation in performance—whether from data drift or changing conditions—is detected and addressed. By managing the full lifecycle of models in production, this layer connects the data architecture to the institution's operational decisions and products, closing the loop between data and action.

6.6 The governance layer

The governance layer spans all the others and, the analysis finds, is the most decisive element of an AI-first architecture. It manages data quality, lineage, access control, privacy, and documentation. In an AI-first institution, governance cannot be an afterthought because any errors, gaps, or biases in the data propagate directly into automated decisions that may affect large numbers of people. Strong governance ensures that data is trustworthy, that its origins and transformations can be traced, that access is appropriately controlled, and that the institution can explain and audit how its AI systems reach their conclusions—an increasingly important requirement in regulated sectors. The governance layer is therefore what distinguishes an architecture that is merely fast from one that is both fast and reliable.

6.7 The central finding: data as a product

A central finding of the analysis is that AI-first architecture requires a fundamental reorientation of the institution's relationship with data. Rather than storing data as a static record of past events, institutions must treat it as a continuously flowing, reusable product. This shift requires not only the technical components described above but also clear ownership, defined quality standards, and a culture in which teams are accountable for the data they produce and share. The analysis suggests that institutions that adopt this data-as-a-product mindset, supported by strong governance, move from experiment to reliable production far more quickly and with fewer failures than those that focus only on acquiring tools and storage.




Figure 3. Traditional architecture versus AI-first architecture.

7. Discussion

7.1 Interpretation of the findings

The proposed reference model demonstrates that designing for AI is not simply a matter of adding more storage or compute to an existing architecture. The distinguishing features of AI-first architecture are reproducibility, timeliness, and reuse, each addressed by specific layers of the model. Reproducibility, enabled by the feature store and by lineage tracking in the governance layer, ensures that models can be trained, audited, and retrained reliably. Timeliness, enabled by streaming ingestion, allows AI to act on current rather than historical data. Reuse, enabled by treating data as a product with clear ownership, prevents the duplication and inconsistency that slow large organisations and undermine trust in their data.

7.2 Theoretical and practical implications

These findings reinforce and extend the view in the literature that the hardest problems in AI are data-related and organisational rather than algorithmic (Sculley et al., 2015). The governance layer emerges as particularly decisive: without it, faster pipelines simply deliver unreliable data more quickly, and automated decisions inherit hidden errors and biases. This has significant practical implications, especially for institutions in regulated sectors such as finance and healthcare, where the ability to explain and audit data flows and model decisions

is not optional but a legal and ethical requirement. For such institutions, investing in governance and lineage is not a cost to be minimised but a prerequisite for deploying AI at all.

7.3 Trade-offs and contextual factors

The analysis also highlights important trade-offs. Decentralising data ownership, as advocated by the data-mesh literature, improves scalability and domain relevance but can fragment standards if shared governance is weak. Real-time streaming pipelines add complexity and cost that not every application justifies; many use cases are served perfectly well by batch processing. The right architecture therefore depends on the institution's scale, sector, and maturity, and the reference model should be adapted rather than copied wholesale. A small organisation may implement a simplified version with a single team and modest streaming capability, while a large, complex institution may require full decentralisation and extensive governance automation. The model is intended as a structured starting point for these decisions, not a prescription.

7.4 Limitations and future research

The principal limitation of this study is that the reference model is evaluated conceptually rather than through implementation and measurement. While it is grounded in established scholarship and documented practice, its benefits have not been quantified in a specific institutional setting. Future research should test the model empirically, measuring outcomes such as time-to-deployment, model reliability, the frequency of data-quality incidents, and the auditability of decisions in institutions that adopt it. Comparative case studies across sectors would also help clarify how the model should be adapted to different regulatory and organisational contexts, and longitudinal studies could examine how AI-first architectures evolve as institutions mature.

8. Conclusion

AI-first institutions depend on an enterprise data architecture designed for machine learning from the outset rather than adapted from systems built for reporting. This paper has argued that such an architecture consists of five layers—ingestion, storage, feature processing, model serving, and a cross-cutting governance layer—and that its defining qualities are reproducibility, timeliness, and reuse. Above all, it has shown that the decisive shift is conceptual as well as technical: data must be treated as a governed, reusable product rather than a static record, supported by clear ownership, quality standards, and an organisational

10.48047/jocaaa.2024.32.01.105

culture that values data. Technology alone is insufficient. Institutions that align their architecture with strong governance and data-as-a-product thinking position themselves to deploy AI reliably and at scale, shortening the path from raw data to deployed model. Those that invest only in storage and tools while neglecting governance risk building systems that are fast but fragile, delivering unreliable data and flawed automated decisions more quickly. The reference model and design principles offered here provide a practical foundation that institutions can adapt to their own scale, sector, and maturity as they build the data foundations for an AI-first future.

References1: 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>

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)

10.48047/jocaaa.2024.32.01.105

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>

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 callyspongiolide; 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>

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>