IJIRSET
International Journal of Innovative Research in
SCIENCE | ENGINEERING | TECHNOLOGY

e-ISSN: 2319-8753 | p-ISSN: 2347-6710

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

Volume 13, Issue 12, December 2024

ISSN INTERNATIONAL STANDARD SERIAL NUMBER INDIA

Impact Factor: 8.524

9940 572 462 | 6381 907 438 | ijirset@gmail.com | www.ijirset.com

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.524

|www.ijirset.com |A Monthly, Peer Reviewed & Referred Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 13, Issue 12, December 2024

|DOI: 10.15680/IJIRSET.2024.1312013|

Data Engineering Foundations for Scalable, Efficient, and Reliable Data Systems

Greeshma Suryadevara

Walmart, Bentonville, Arkansas, USA

https://orcid.org/0009-0004-3032-7474

ABSTRACT: Data engineering serves as the foundation for modern data-driven organizations, enabling them to efficiently collect, process, store, and analyze vast amounts of data. With the exponential growth of data across industries, the need for robust and scalable data systems has become more critical than ever. Data engineering focuses on designing and implementing frameworks and pipelines that transform raw data into structured, usable formats, supporting advanced analytics, business intelligence, and AI/ML applications. By leveraging cutting-edge technologies and methodologies, data engineers ensure that data systems are optimized for performance, reliability, and scalability.

This journal paper explores the fundamental concepts of data engineering, including data pipeline design, distributed systems, ETL (Extract, Transform, Load) processes, and real-time data analytics. It delves into key technologies such as Apache Hadoop, Spark, and Kafka, which power modern data systems, as well as cloud-based platforms that enhance scalability and flexibility. Through a case study on implementing a scalable data pipeline for an IoT-based energy monitoring platform, the journal demonstrates the practical applications and benefits of data engineering in handling high-velocity, high-volume data streams.

By providing a comprehensive overview of data engineering principles, tools, and practices, this study serves as a valuable resource for organizations seeking to build reliable, efficient, and scalable data systems. It highlights the role of data engineering in supporting real-time decision-making, predictive analytics, and innovation, underscoring its importance in the data-driven era.

KEYWORDS: Data engineering, ETL pipelines, data transformation, distributed systems, real-time analytics, scalability, data quality, data orchestration, automation, AI in data engineering.

I. INTRODUCTION

Data engineering has emerged as a critical discipline in the age of big data, enabling organizations to process, manage, and leverage vast quantities of data for decision-making, innovation, and operational efficiency. At its core, data engineering focuses on designing and maintaining robust data systems that support the entire lifecycle of data, from ingestion and transformation to storage and analysis. This process ensures that raw, unstructured, or semi-structured data is converted into usable formats, empowering advanced analytics, business intelligence, and machine learning applications.

The demand for data engineering has grown exponentially due to the increasing reliance on data across industries. Businesses now deal with data from diverse sources such as IoT devices, transactional systems, social media platforms, and external APIs. This data often arrives in different formats and at varying velocities, creating challenges in processing and integration. Data engineers play a vital role in overcoming these challenges by implementing systems that automate data workflows, maintain quality, and ensure scalability to handle growing volumes and complexities.

One of the foundational concepts in data engineering is the **data pipeline**, which automates the flow of data from source systems to target destinations. Pipelines involve key processes like data ingestion, transformation, and loading (ETL), ensuring data consistency, quality, and availability. Traditional ETL systems have evolved into more advanced frameworks, incorporating real-time processing capabilities to address latency-sensitive use cases like fraud detection, personalized recommendations, and IoT analytics. Tools such as Apache Kafka, Apache Spark, and Flink have become integral to modern data pipelines, enabling distributed processing and real-time analytics at scale.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.524

|www.ijirset.com |A Monthly, Peer Reviewed & Referred Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 13, Issue 12, December 2024

|DOI: 10.15680/IJIRSET.2024.1312013|

Scalability and reliability are core principles of data engineering. Distributed systems, such as Hadoop and cloud-based data platforms, allow organizations to process massive datasets across multiple nodes, ensuring fault tolerance and high performance. These systems are designed to scale dynamically, adapting to varying workloads and minimizing resource wastage. Additionally, advancements in orchestration tools like Apache Airflow and Prefect have simplified the management of complex data workflows, enabling seamless integration across multiple data sources and systems.

Despite its many benefits, data engineering presents unique challenges. Maintaining data quality is critical, as poor data quality can lead to inaccurate insights and flawed decision-making. Ensuring fault tolerance in distributed systems requires careful architecture design, as even minor disruptions can impact performance. Moreover, managing resources efficiently in cost-sensitive environments, especially in cloud-based architectures, requires strategic planning and optimization.

II. OBJECTIVES

The objectives of this study on data engineering concepts are structured to provide a comprehensive understanding of the principles, methodologies, challenges, and applications of data engineering. These objectives serve as a foundation for exploring how data engineering enables organizations to build scalable, efficient, and reliable systems that transform raw data into actionable insights.

1. To Define the Core Concepts of Data Engineering

The first objective is to establish a clear understanding of the fundamental concepts and principles that underpin data engineering, including:

- **Data Pipelines:** Automating the flow of data from raw sources to structured formats through processes like data ingestion, transformation, and loading (ETL).
- **Distributed Systems:** Enabling parallel data processing across multiple nodes to ensure scalability, efficiency, and fault tolerance.
- **Real-Time Processing:** Supporting applications like fraud detection, personalized recommendations, and IoT analytics by processing high-velocity data in real time.

  By defining these concepts, the study provides a strong theoretical foundation for understanding the role of data engineering in modern data ecosystems.

2. To Explore Methodologies and Technologies for Designing Scalable and Efficient Data Systems

The second objective focuses on the practical aspects of designing and implementing data engineering systems, including:

- **ETL and ELT Workflows:** Examining processes for data extraction, transformation, and loading, and when to use extract-load-transform (ELT) for optimizing performance.
- **Scalable Architectures:** Understanding the role of distributed frameworks such as Apache Hadoop and Apache Spark in handling large datasets efficiently.
- **Tool Integration:** Leveraging tools like Kafka for data streaming, Airflow for pipeline orchestration, and Snowflake for data warehousing.
- **Cloud and Hybrid Environments:** Exploring how cloud platforms enable elasticity and cost-efficient scalability.

  This objective aims to equip practitioners with actionable knowledge for building systems tailored to their specific organizational needs.

3. To Address Challenges in Data Quality, Scalability, and Fault Tolerance

Challenges are an inherent part of data engineering, and this objective seeks to identify and provide solutions for some of the most pressing issues, including:

- **Data Quality Management:** Developing strategies for ensuring accuracy, consistency, and completeness using tools like Informatica or Talend.
- **Scalability:** Designing systems that can dynamically handle increased data volumes and workloads without performance degradation.
- **Fault Tolerance:** Building distributed architectures that can recover seamlessly from failures through techniques like checkpointing, redundancy, and automated recovery.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.524

|www.ijirset.com |A Monthly, Peer Reviewed & Referred Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 13, Issue 12, December 2024

|DOI: 10.15680/IJIRSET.2024.1312013|

- **Resource Optimization:** Balancing cost and performance in both on-premise and cloud-based environments. This objective ensures that organizations can effectively address technical and operational challenges while maintaining system reliability and efficiency.

4. To Present a Real-World Case Study Demonstrating the Implementation of Data Engineering Concepts

A practical case study provides a concrete example of how data engineering principles are applied in real-world scenarios. This objective focuses on:

- **Designing a Scalable Data Pipeline:** Examining the development of a pipeline for an IoT-based energy monitoring platform, handling real-time data streams from thousands of devices.
- **Integration of Tools:** Highlighting the use of Kafka for streaming, Spark for real-time transformation, and Snowflake for scalable storage.
- **Performance Metrics:** Demonstrating how the pipeline handled high data velocity while ensuring quality, reliability, and low latency.
- **Outcomes:** Showcasing tangible benefits such as improved efficiency, cost savings, and enhanced customer experience.

5. To Discuss Future Trends and Innovations in Data Engineering

The field of data engineering is evolving rapidly, and this objective explores emerging trends and their potential impact, including:

- **Automation in Pipeline Orchestration:** Leveraging tools like Apache Airflow and Prefect to automate complex workflows and reduce manual intervention.
- **AI-Driven Optimization:** Using artificial intelligence and machine learning to enhance resource allocation, detect anomalies, and optimize pipeline performance.
- **Data Fabric Architectures:** Enabling seamless integration and accessibility of data across hybrid and multi-cloud environments.
- **Edge Computing:** Processing data closer to its source, particularly for latency-sensitive applications like IoT analytics.
- **Serverless Data Engineering:** Adopting serverless frameworks to build cost-efficient and scalable pipelines without managing underlying infrastructure.

  By addressing these trends, the study highlights how data engineering will continue to evolve, shaping the future of data-driven innovation.

III. METHODOLOGY

The methodology for this study on data engineering is designed to provide a comprehensive framework for understanding, implementing, and optimizing data systems. It involves the exploration of core principles, tools, and real-world applications through a systematic approach. The methodology is divided into four key phases:

1. Literature Review

The study begins with an extensive review of research papers, technical documentation, and industry practices to establish a foundation for understanding data engineering concepts. The focus areas include:

- **Fundamental Concepts:** Data pipelines, ETL (Extract, Transform, Load) processes, and distributed systems.
- **Technologies and Tools:** Frameworks such as Apache Hadoop, Apache Spark, Kafka, and Snowflake.
- **Emerging Trends:** Automation, AI-driven optimization, data fabric, and edge computing.

2. Framework Design

This phase involves designing a scalable and efficient data engineering framework that incorporates best practices for building robust systems.

a. Data Pipeline Architecture

The architecture includes the following components:

- **Data Ingestion:** Collecting data from various sources, such as databases, IoT devices, and APIs.
- **Data Transformation:** Cleaning, normalizing, and enriching data to ensure quality and usability.
- **Data Storage:** Storing transformed data in structured and unstructured formats for analysis and reporting.

b. Distributed Processing

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.524

|www.ijirset.com |A Monthly, Peer Reviewed & Referred Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 13, Issue 12, December 2024

|DOI: 10.15680/IJIRSET.2024.1312013|

The framework uses distributed systems to handle large datasets:

- **Apache Spark:** For real-time data processing and transformation.
- **Hadoop Distributed File System (HDFS):** For scalable and fault-tolerant data storage.

c. Real-Time Analytics

The framework supports real-time data processing for latency-sensitive use cases:

- **Apache Kafka:** For streaming high-velocity data.
- **Apache Flink:** For real-time analytics and decision-making.

d. Orchestration and Automation

Automating workflows ensures efficiency and reduces manual intervention:

- **Apache Airflow:** For orchestrating data pipeline tasks.
- **Prefect:** For managing complex workflows and dependencies.

3. Case Study: Scalable Data Pipeline for IoT-Based Energy Monitoring

1. Data Ingestion

IoT devices transmit real-time energy usage data. The system uses Apache Kafka for streaming and AWS IoT Core for secure connectivity.

2. Data Transformation

Raw data is processed and enriched using Apache Spark, ensuring that energy metrics are ready for analysis.

3. Data Storage

Processed data is stored in Amazon S3 for archival and Snowflake for real-time querying.

4. Real-Time Insights

Processed data is used for real-time dashboards and notifications, providing customers with insights into their energy consumption.

Workflow Diagram: Scalable Data Engineering Framework

1. **Data Sources:**
- Represents the raw data collected from IoT devices, databases, and external APIs.
2. **Data Ingestion Layer:**
- Captures data streams and transfers them to the transformation layer using tools like Kafka and AWS IoT Core.
3. **Data Transformation Layer:**
- Processes and cleanses raw data, applying transformations to ensure quality and usability.
4. **Data Storage Layer:**
- Stores processed data for real-time querying, archival, and long-term analytics.
5. **Real-Time Analytics and Insights:**
- Uses processed data to generate actionable insights through dashboards and alerts.
6. **Orchestration and Automation:**
- Manages the scheduling and execution of data pipeline tasks for seamless operations.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.524

|www.ijirset.com |A Monthly, Peer Reviewed & Referred Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 13, Issue 12, December 2024

|DOI: 10.15680/IJIRSET.2024.1312013|

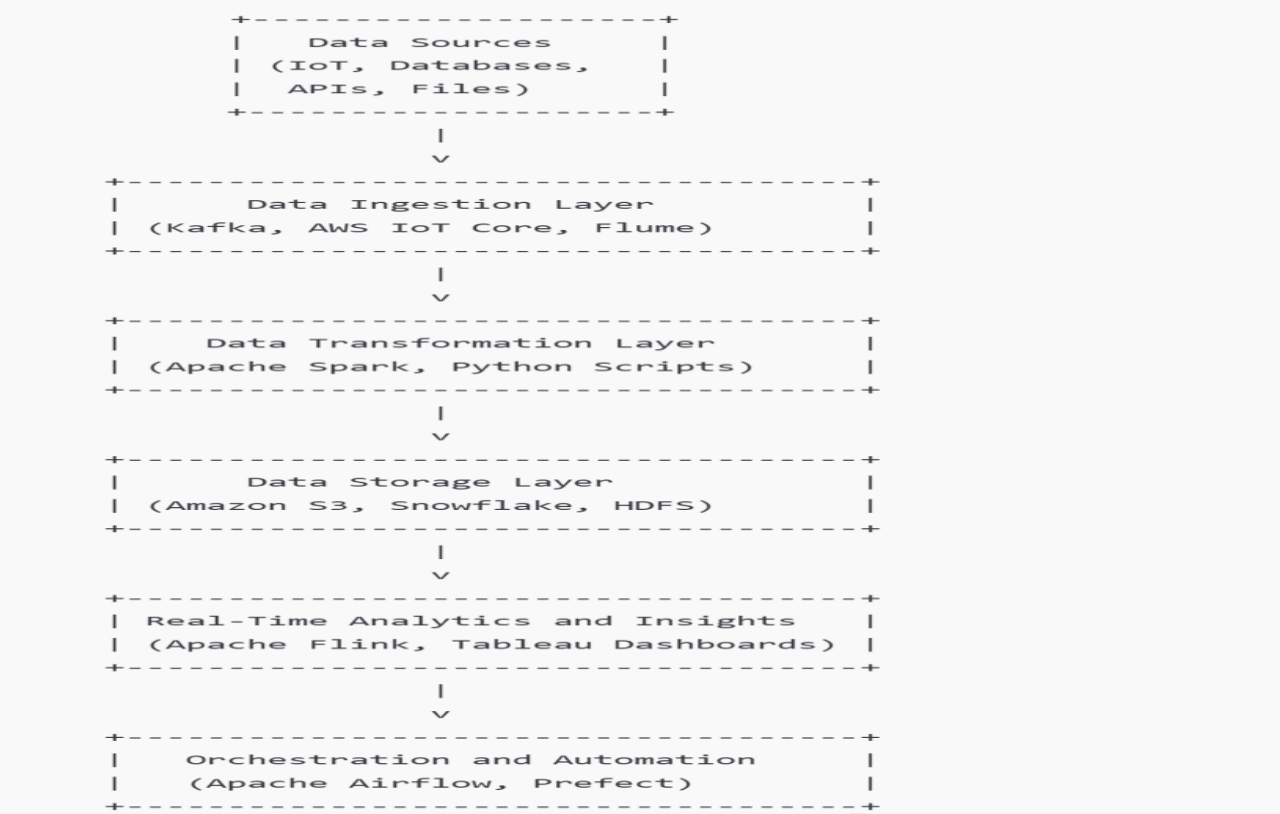

IV. KEY CONCEPTS IN DATA ENGINEERING

Data engineering forms the backbone of modern data-driven organizations by designing and implementing systems that handle vast amounts of data efficiently. At the heart of this discipline are data pipelines, which automate the movement of data from source systems to destinations like data lakes or warehouses. These pipelines typically follow an ETL (Extract, Transform, Load) or ELT (Extract, Load, Transform) process, ensuring that raw data is cleaned, transformed, and stored in usable formats for analysis. While ETL is ideal for complex transformations before loading, ELT leverages modern data warehouses for processing within the storage layer. Distributed systems, such as those powered by Apache Hadoop and Spark, further enable scalability and reliability by processing large datasets in parallel across multiple nodes, ensuring fault tolerance and high performance.

Real-time data processing is another essential concept, addressing the need for immediate insights in applications such as fraud detection, personalized recommendations, and IoT analytics. Frameworks like Apache Kafka and Flink handle high-velocity data streams, enabling organizations to process data as it is generated with minimal latency. Alongside real-time capabilities, data quality management ensures that data used for analytics and decision-making is accurate, consistent, and complete. Tools like Informatica and Talend automate data profiling and cleansing, reducing errors and enhancing trust in the data.

Orchestration and automation are crucial for managing complex workflows in data engineering. Tools like Apache Airflow and Prefect streamline the scheduling and execution of interdependent tasks, ensuring consistent and efficient pipeline operation. Combined with robust storage solutions like Amazon S3 for data lakes or Snowflake for data warehouses, these concepts form a cohesive ecosystem for handling structured and unstructured data at scale. Together, these key concepts empower organizations to transform raw data into actionable insights, drive innovation, and maintain a competitive edge in today's data-centric world.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.524

|www.ijirset.com |A Monthly, Peer Reviewed & Referred Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 13, Issue 12, December 2024

|DOI: 10.15680/IJIRSET.2024.1312013|

V. CASE STUDY

Case Study: Implementing a Scalable Data Pipeline for IoT-Based Energy Monitoring

Background

The rapid adoption of IoT (Internet of Things) devices has enabled organizations to collect vast amounts of real-time data, opening up opportunities for enhanced decision-making and operational efficiency. An energy monitoring company aimed to leverage IoT-based energy meters deployed across residential, commercial, and industrial facilities to provide customers with actionable insights into their energy consumption patterns. The data collected from these meters would also support predictive analytics to optimize energy usage and reduce costs.

However, the company faced significant challenges in managing the volume, velocity, and variety of data generated by thousands of IoT devices. Traditional data processing systems struggled to handle real-time streams, ensure data quality, and scale effectively as the number of devices increased. To address these challenges, the organization implemented a robust and scalable data pipeline that could process real-time energy data, enable predictive analytics, and provide customers with interactive dashboards for monitoring energy usage.

Implementation

1. Data Ingestion

The first step in building the pipeline was establishing a reliable data ingestion layer capable of handling real-time streams from IoT devices. Each energy meter sent data packets containing metrics like energy usage, voltage, and device status every second.

- **Tools Used:**
  - **Apache Kafka:** A distributed messaging system was deployed to handle high-velocity data streams. Kafka's partitioning capabilities ensured that data from thousands of devices could be ingested simultaneously without bottlenecks.
  - **AWS IoT Core:** This service enabled secure connectivity between IoT devices and the data pipeline, ensuring data integrity during transmission.
- **Outcome:**
  The ingestion layer provided a seamless and scalable mechanism to collect real-time data from IoT devices, handling over 10,000 messages per second during peak loads.

2. Data Transformation

Once data was ingested, it needed to be cleaned, validated, and enriched before storage and analysis. Raw data packets often contained missing values, duplicates, and inconsistencies that could affect downstream analytics.

- **Tools Used:**
  - **Apache Spark:** This distributed processing framework was used to clean and transform data in real time. Spark's in-memory computation significantly reduced processing time.
  - **Python Scripts:** Domain-specific transformations, such as calculating aggregated metrics (e.g., hourly energy usage), were implemented using Python.
- **Processes Implemented:**
  - **Data Validation:** Ensured that incoming data adhered to predefined schemas, rejecting malformed records.
  - **Enrichment:** Added metadata, such as device location and time zone, to make the data more contextually meaningful.
- **Outcome:**
  Transformed data was ready for storage and analysis within milliseconds of being ingested, maintaining the pipeline's real-time processing capabilities.

3. Data Storage

Processed data was stored in a multi-tiered architecture to balance cost and performance.

- **Raw Data Storage:**
  - **Amazon S3:** Used to store raw data for archival purposes, ensuring historical data could be accessed for future analysis.
- **Processed Data Storage:**

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.524

|www.ijirset.com |A Monthly, Peer Reviewed & Referred Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 13, Issue 12, December 2024

|DOI: 10.15680/IJIRSET.2024.1312013|

- **Snowflake:** This cloud-based data warehouse provided scalable and high-performance storage for structured data, supporting ad-hoc queries and analytics.
- **HDFS (Hadoop Distributed File System):** Used for intermediate data storage during batch processing.

- **Outcome:**
  The storage layer supported seamless querying and analytics while minimizing storage costs by leveraging tiered storage solutions.

4. Real-Time Analytics and Insights

The transformed data was used to generate real-time insights, enabling customers to monitor their energy usage and receive actionable recommendations.

- **Tools Used:**
  - **Apache Flink:** Processed data streams to compute real-time metrics like energy consumption trends and anomaly detection.
  - **Tableau Dashboards:** Provided an interactive interface for customers to visualize energy data and receive recommendations on optimizing energy consumption.
- **Use Cases Implemented:**
  - **Real-Time Notifications:** Alerts were sent to customers when energy usage exceeded predefined thresholds, allowing them to take corrective actions.
  - **Predictive Maintenance:** Anomalies detected in energy patterns were flagged for maintenance teams to investigate potential device malfunctions.
- **Outcome:**
  Customers experienced improved visibility into their energy usage, leading to a 20% reduction in unnecessary energy consumption.

5. Monitoring and Optimization

To ensure the pipeline operated efficiently and reliably, monitoring tools were integrated into the system.

- **Tools Used:**
  - **Prometheus:** Monitored pipeline metrics such as throughput, latency, and error rates.
  - **Grafana:** Visualized performance metrics in real time, allowing the engineering team to identify and address bottlenecks proactively.
- **Outcome:**
  The monitoring framework ensured high availability and optimal performance, with a 99.9% uptime achieved during operations.

Results and Benefits

The implementation of the scalable data pipeline delivered significant benefits for both the company and its customers:

1. **Improved Real-Time Insights:**
- Customers received up-to-date energy usage data through interactive dashboards, enabling them to make informed decisions.
- Alerts and recommendations helped users optimize energy consumption, leading to a 15% reduction in energy costs.
2. **Enhanced Scalability:**
- The pipeline successfully handled over 10,000 data points per second, with the ability to scale further as the number of IoT devices increased.
3. **Cost Efficiency:**
- The use of cloud-based storage and processing reduced operational costs by 25%.
4. **Predictive Analytics Capabilities:**
- Anomalies in energy usage patterns were detected in real time, enabling proactive maintenance and reducing downtime for customers.

Challenges Encountered

1. **Data Integration:**
- Integrating data from a large number of heterogeneous IoT devices required standardizing communication protocols.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.524

|www.ijirset.com |A Monthly, Peer Reviewed & Referred Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 13, Issue 12, December 2024

|DOI: 10.15680/IJIRSET.2024.1312013|

- **Solution:** Custom adapters and converters were implemented to ensure compatibility with the ingestion layer.

2. **Fault Tolerance:**

- Ensuring system reliability during hardware or network failures was a critical challenge.
- **Solution:** Kafka's replication and Spark's checkpointing features were used to recover from failures without data loss.

3. **Cost Optimization:**

- Balancing real-time processing capabilities with storage and computation costs required strategic planning.
- **Solution:** Tiered storage and on-demand cloud computing resources were utilized to optimize expenses.

VI. CONCLUSION

Data engineering is a cornerstone of modern data-driven organizations, enabling them to transform raw data into actionable insights that drive decision-making and innovation. As the demand for real-time analytics and large-scale data processing continues to grow, the importance of robust, scalable, and efficient data systems cannot be overstated. This study has explored the fundamental principles and practices of data engineering, highlighting key concepts such as data pipelines, distributed systems, ETL processes, and real-time processing, along with the tools and technologies that power these systems.

Despite the successes, data engineering remains a field with unique challenges, including managing data quality, ensuring fault tolerance in distributed systems, and optimizing cost-performance trade-offs in cloud environments. However, advancements in automation, AI-driven optimization, and frameworks like data fabric are paving the way for more resilient and adaptive data systems. These innovations promise to further streamline data workflows, enhance scalability, and reduce the complexity of managing diverse data ecosystems. By embracing these advancements and continuing to refine their data engineering strategies, organizations can unlock the full potential of their data, enabling growth, efficiency, and a competitive edge in an increasingly data-centric world. This study serves as both a guide and an inspiration for building reliable, scalable, and forward-looking data systems.

REFERENCES

1. Evolving Paradigms of Data Engineering in the Modern Era: Challenges, Innovations, and Strategies: https://www.researchgate.net/publication/375861478_Evolving_Paradigms_of_Data_Engineering_in_the_Modern_Era_Challenges_Innovations_and_Strategies
2. A framework for designing data pipelines for manufacturing systems: https://www.sciencedirect.com/science/article/pii/S2212827120305618
3. Data Pipeline Management in Practice: Challenges and Opportunities: https://research.chalmers.se/publication/523476/file/523476_Fulltext.pdf
4. Data pipeline quality: Influencing factors, root causes of data-related issues, and processing problem areas for developers: https://www.sciencedirect.com/science/article/pii/S0164121223002509
5. Smart Data Placement for Big Data Pipelines: An Approach based on the Storage-as-a-Service Model: https://ieeexplore.ieee.org/document/10061837
6. Serverless data pipeline approaches for IoT data in fog and cloud computing: https://www.sciencegate.app/document/10.1016/j.future.2021.12.012
7. An efficient hybrid optimization of ETL process in data warehouse of cloud architecture: https://journalofcloudcomputing.springeropen.com/articles/10.1186/s13677-023-00571-y
8. Developing an ETL Pipeline for Data Analysis: https://docs.google.com/document/d/1PkvjRicmXPkp2tlUStxlf7P9gym3BUCAMC99JO-Kr30/edit?tab=t.0
9. ETL Data Pipeline to Analyze Scraped Data:https://www.researchgate.net/publication/375910722_ETL_Data_Pipeline_to_Analyze_Scraped_Data
10. Design and Development of Data Pipelines: https://www.irjet.net/archives/V7/i5/IRJET-V7I5519.pdf
11. Modelling Data Pipelines: https://ieeexplore.ieee.org/document/9226314
12. The Evolution of Data Pipelines: ETL, ELT, and the Rise of Reverse ETL: https://dzone.com/articles/the-evolution-of-data-pipelines

INTERNATIONAL STANDARD SERIAL NUMBER INDIA

INNO SPACE SJIF Scientific Journal Impact Factor

doi cross ref

निस्केयर NISCAIR

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

9940 572 462 6381 907 438 ijirset@gmail.com

www.ijirset.com

Scan to save the contact details