Data Classification: Enabling Robust Data Governance and Access Management in Enterprise Environments

Dinesh Thangaraju

AWS Data Platform
Amazon Web Services, Amazon.com Corp LLC
Seattle, United States of America
thangd@amazon.com

Abstract

In the era of data-driven decision-making, enterprises face challenges in managing and protecting their data assets amidst regulatory compliance requirements and increasing cybersecurity threats. Data classification is the foundation for enabling robust access management, ensuring data quality, enforcing consistent access control policies, and strengthening data governance. This paper highlights the importance of data classification in modern enterprises, outlines key challenges, and proposes a technical framework for implementing a scalable data classification solution. By integrating automation, metadata management, and policy enforcement mechanisms, businesses can secure sensitive data, improve operational efficiency, and ensure compliance with global regulations. Metrics for evaluating implementation success are also presented to help enterprises achieve a resilient data management infrastructure.

Keywords: Data Classification, Data Governance, Access Control, Metadata Management, Compliance, Security Policies, Data Quality

I. INTRODUCTION

In the era of big data and digital transformation, organizations face unprecedented challenges in managing, securing, and deriving value from their vast data assets. The exponential growth of enterprise data has transformed information into a strategic asset. However, this transformation has introduced challenges related to **security**, **compliance**, and **governance**. Organizations must not only protect sensitive data but also ensure appropriate **access controls**, **policy enforcement**, and **data quality standards**.

Data classification emerges as a critical component in addressing these challenges, providing a structured approach to categorizing data based on its sensitivity, value, and regulatory requirements. This paper explores how data classification enables robust access management, ensures data quality, and supports consistent policy-based access control, ultimately contributing to strong data governance within the enterprise.

The proliferation of data sources, coupled with increasing regulatory pressures and security threats, necessitates a systematic approach to data management. Data classification serves as the cornerstone of this approach, allowing organizations to:

- Implement granular access controls based on data sensitivity

- Prioritize data quality efforts for critical datasets
- Ensure compliance with regulatory requirements
- Optimize data storage and processing resources
- Enhance overall data security posture

The Need for Data Classification

Unstructured and fragmented data poses security risks, reduces operational efficiency, and leads to compliance gaps. For instance, without proper classification, organizations may expose personally identifiable information (PII) to unauthorized users, leading to data breaches and regulatory fines.

Key Questions Addressed in this Paper:

1.What challenges do enterprises face without a data classification framework?

2.What technical approaches can be adopted to implement scalable data classification?

3.How does data classification enable effective governance, access management, and compliance?

This paper provides a technical roadmap for deploying data classification solutions, detailing challenges, design considerations, and evaluation metrics to ensure scalability and compliance.

II. Challenges in Enterprise Data Management

A. Data Volume and Diversity

Modern enterprises often deal with vast amounts of data from a wide range of sources, such as customer transactions, IoT sensor data, social media interactions, and enterprise applications. This explosion of data volume and diversity makes it challenging for organizations to maintain a consistent and scalable data classification system across their entire data landscape. For example, a large retail company may have customer data stored in their e-commerce platform, point-of-sale systems, and customer relationship management (CRM) tools. Each of these data sources could have its own classification schema, making it difficult to apply uniform security controls, ensure data quality, and enforce consistent access policies enterprise-wide.

The lack of a centralized and consistent data classification approach can have significant consequences for the organization Without a clear understanding of data sensitivity and criticality, enterprises risk exposing sensitive information like personally identifiable data to unauthorized users, leading to potential data breaches and regulatory fines. Additionally, the inability to prioritize data management efforts based on classification can result in inefficient resource allocation, poor data quality, and compliance gaps. It also leads to a lack of visibility into data lineage and ownership, as well as difficulty in enforcing policies uniformly across heterogeneous systems.

B. Evolving Regulatory Landscape

Compliance requirements such as the General Data Protection Regulation (GDPR), the California Consumer Privacy Act (CCPA), and industry-specific regulations necessitate dynamic data classification systems that can adapt to changing legal frameworks. As new privacy laws and data protection mandates emerge globally, organizations must have the ability to quickly identify and classify sensitive data elements, such as personally identifiable information (PII), to ensure compliance.

The inability to keep pace with the evolving regulatory landscape can have severe consequences for enterprises. Failure to properly classify and protect sensitive data can lead to costly data breaches, hefty

fines, and reputational damage. For example, a GDPR violation can result in penalties of up to 4% of a company's global annual revenue or €20 million, whichever is higher. Similarly, non-compliance with CCPA can incur civil penalties of up to $7,500 per violation.

C. Siloed Data Environments, Legacy Systems and Data

Many organizations struggle with data silos, where different departments or business units maintain separate classification schemes, hindering enterprise-wide governance efforts. This fragmented approach to data classification makes it extremely difficult to apply uniform security controls, ensure data quality, and enforce consistent access policies across the organization.

Additionally, integrating legacy systems and historical data into modern classification frameworks poses significant technical and operational challenges. Legacy applications and databases often lack the metadata and contextual information required for accurate classification, leading to inconsistencies and gaps in the overall data governance strategy. The inability to seamlessly incorporate legacy data into a centralized classification system results in incomplete visibility, inefficient resource allocation, and compliance risks.

The lack of a unified, enterprise-wide data classification system creates significant challenges for organizations. Without a holistic view of their data assets and their associated sensitivity levels, enterprises risk exposing sensitive information to unauthorized users, leading to potential data breaches and regulatory fines. It also hinders their ability to prioritize data management efforts, ensure data quality, and demonstrate compliance with evolving privacy laws and industry regulations. Furthermore, the siloed nature of classification schemes and legacy data integration issues make it difficult to enforce policies uniformly across the organization, leading to inconsistent access controls and data governance practices.

D. Data Security and Unauthorized Access

Without proper data classification, sensitive information can be exposed to unauthorized users, increasing the risk of data breaches. When enterprises lack a structured approach to categorizing their data assets based on sensitivity levels, it becomes challenging to apply the appropriate security controls and access management policies.

The lack of a comprehensive data classification system can have severe consequences for enterprises. Weak access management leads to insider threats, where employees or malicious actors may gain unauthorized access to critical or regulated data, putting the organization at risk. Additionally, the difficulties in tracking data sharing and movement hinder the ability to detect anomalies, investigate potential data leaks, and enforce data handling policies consistently. Without visibility into how sensitive information is being accessed and distributed, enterprises become more vulnerable to data breaches, which can result in financial losses, regulatory fines, and reputational damage.

E. Scalability and Automation

As data volumes continue to grow exponentially, manual classification methods become increasingly ineffective and unsustainable. Relying on manual processes to categorize and label data assets across the enterprise leads to inconsistent categorization, delayed enforcement, and an inability to adapt to real-time classification needs.

The lack of scalable, automated data classification methods can have significant consequences for enterprises. Without a centralized, automated system, organizations struggle to keep pace with the rapid

influx of new data and changing business requirements, leading to an inability to adapt to real-time classification needs. This, in turn, leaves sensitive data vulnerable to exposure as the appropriate security controls and access policies cannot be applied in a timely manner. Furthermore, the fragmented and inconsistent policies that arise from manual classification make it challenging to enforce policies uniformly, track data access and usage, and maintain comprehensive audit trails for compliance and security purposes.

IV. Technical Approach to Data Classification

A. Architecture Overview

A modern data classification framework integrates several key components to enable automated and scalable data classification workflows:

1. Metadata Repository: At the core of the framework is a centralized database for storing classification labels, sensitivity levels, and ownership tags associated with the data assets. This metadata repository serves as the single source of truth for the organization's data classification schema.
2. Machine Learning Models: The framework leverages advanced machine learning algorithms to automatically detect and classify sensitive data elements, such as personally identifiable information (PII), financial data, and other regulated content. By analyzing the content, context, and patterns within the data, these ML models can accurately categorize data assets without relying solely on manual processes.
3. Policy Engine: The data classification framework integrates with a policy engine that enforces access controls and compliance policies based on the assigned classification labels. This allows the organization to implement granular, role-based or attribute-based access management, ensuring that users can only access data appropriate to their level of authorization.
4. Audit Logs and Reporting: To provide transparency and traceability, the framework includes comprehensive audit logging and reporting capabilities. This enables the organization to monitor data access patterns, investigate potential policy violations, and demonstrate compliance with regulatory requirements.

By integrating these key components - metadata management, machine learning, and policy enforcement - the modern data classification framework empowers enterprises to automate the categorization of their data assets, apply appropriate security controls, and maintain a robust data governance infrastructure.

B. Workflow for Implementation

To address the challenges posed by growing data volumes, evolving regulatory landscapes, and siloed data environments, enterprises must adopt a comprehensive workflow for deploying a modern data classification framework.

The first step is Data Discovery and Profiling, where organizations leverage data cataloging tools to scan their data sources and gather detailed metadata. This is crucial in overcoming the challenges of data silos, as it provides a holistic view of the enterprise's data assets, regardless of their origin or storage location. By using machine learning-based algorithms to identify and tag sensitive data elements, such as personally identifiable information (PII), enterprises can mitigate the risks of unauthorized access and data breaches.

Next, the organization must define Classification Labels and Taxonomy that align with their data governance policies and regulatory requirements. Establishing a clear hierarchy of classification levels, from Public to Highly Sensitive, enables enterprises to apply the appropriate security controls and access

management policies. This is particularly important in addressing the evolving regulatory landscape, as the classification schema must be dynamic and adaptable to changing compliance mandates like GDPR and CCPA.

The third step involves Policy Definition and Enforcement, where enterprises create role-based (RBAC) and attribute-based (ABAC) access control policies tied to the classification tags. By automating the enforcement of these policies using Identity and Access Management (IAM) tools, organizations can overcome the challenges of fragmented policies and lack of audit trails, ensuring consistent data access and usage across the enterprise.

To further strengthen data security, the framework includes Encryption and Masking capabilities, where sensitive data is protected based on its classification level. This helps address the challenge of data security and unauthorized access, as enterprises can apply granular controls to safeguard critical information.

Finally, the workflow includes Monitoring and Compliance Reporting mechanisms, such as compliance dashboards and periodic data audits. This enables enterprises to maintain visibility into their data classification status, identify anomalies, and demonstrate compliance with regulatory requirements. This is crucial in overcoming the scalability and automation challenges, as the framework provides a centralized, automated approach to data governance and compliance.

By implementing this comprehensive data classification workflow, enterprises can establish a resilient and scalable data management infrastructure that addresses the key challenges they face in the modern, data-driven business landscape.

C. Automated Classification Techniques

To address the scalability and automation challenges associated with manual data classification, the modern data classification framework leverages several advanced techniques:

- Machine Learning Algorithms:
  - The framework utilizes machine learning algorithms to analyze the content and context of data assets for accurate classification.
  - By training these ML models on a diverse dataset of labeled examples, the system can automatically detect and categorize sensitive information, such as personally identifiable data, financial records, or other regulated content.
  - This automated approach to classification allows enterprises to keep pace with the rapid influx of new data, ensuring that appropriate security controls and access policies are applied in a timely manner.
- Pattern Recognition:
  - In addition to machine learning, the framework employs predefined patterns and rules to identify and classify sensitive information.
  - These rules-based techniques can be particularly effective for detecting well-known data types or formats, such as credit card numbers, social security IDs, or other structured data elements.
  - By combining pattern recognition with machine learning, the classification system can leverage the strengths of both approaches to achieve a more comprehensive and accurate categorization of data assets.
- Metadata Analysis:
  - The framework also utilizes file metadata, such as file type, owner, creation date, and other contextual information, to infer appropriate classification levels.

- This metadata-driven approach can be especially useful for classifying unstructured data, like documents or emails, where the content itself may not provide clear indicators of sensitivity.
- By analyzing the metadata associated with these data assets, the classification system can make informed decisions about the appropriate security controls and access policies to apply.

By integrating these automated classification techniques, the modern data classification framework empowers enterprises to overcome the scalability and consistency challenges inherent in manual classification processes. This shift towards a more intelligent, data-driven approach to categorization is essential for organizations to effectively manage and protect their growing data assets in the digital age.

D. Classification Schema Design

When implementing a data classification framework, enterprises must carefully consider the design of the classification schema to ensure it meets their specific needs and constraints.

- Hierarchical vs. Flat Classification Models:
  - Hierarchical classification models organize data assets into a multi-level taxonomy, with broader categories at the top and more granular subcategories below.
  - Flat classification models, on the other hand, use a single level of classification labels without any hierarchical structure.
  - Enterprises must evaluate the pros and cons of these different approaches to determine the most suitable model for their organization. Hierarchical models offer more nuanced control and visibility, but may be more complex to implement and maintain. Flat models are simpler, but may lack the flexibility to address diverse data types and sensitivity levels.
- Granularity Considerations:
  - The level of granularity in the classification schema is another important design decision. Enterprises must balance the need for detailed, precise categorization with the practical constraints of implementation and user adoption.
  - Highly granular classification schemas, with numerous classification levels and subcategories, can provide more accurate and tailored controls. However, they may also increase the complexity of the system, making it more challenging for users to understand and apply the appropriate labels.
  - Conversely, a more coarse-grained classification schema with fewer, broader categories may be easier to implement and use, but may not offer the same level of nuanced control over data access and security.

By carefully considering the trade-offs between hierarchical vs. flat models and the appropriate level of granularity, enterprises can design a classification schema that aligns with their data governance objectives, regulatory requirements, and operational constraints. This thoughtful approach to classification schema design is a crucial step in ensuring the long-term success and scalability of the data classification framework.

E. Integration with Existing Systems

To ensure the seamless adoption and effectiveness of the data classification framework, it must be tightly integrated with the organization's existing data management infrastructure. This includes:

- Data Catalog Integration:
  - Enterprises should ensure that the classification metadata, such as sensitivity levels, ownership tags, and policy associations, are captured and maintained within their enterprise data catalogs.

    - By integrating the data classification framework with the data catalog, organizations can provide a centralized and authoritative source of information about their data assets, including their sensitivity and access control requirements.
    - This integration enables data consumers to easily understand the classification of the data they are accessing, and helps enforce the appropriate security controls and usage policies.
- ETL Process Enhancements:
    - The data classification framework should also be incorporated into the organization's data ingestion and transformation workflows, known as Extract, Transform, and Load (ETL) processes.
    - By embedding the classification logic into these ETL pipelines, enterprises can automatically categorize data as it is ingested from various sources, ensuring that the appropriate labels and metadata are applied from the very beginning.
    - This proactive approach to classification, rather than relying on retroactive labeling, helps maintain data integrity and consistency throughout the data lifecycle, supporting robust data governance and access management initiatives.

By tightly integrating the data classification framework with existing systems, such as data catalogs and ETL processes, enterprises can ensure that classification metadata is captured, maintained, and leveraged across the entire data management ecosystem. This holistic approach helps organizations enforce data policies, maintain data quality, and provide data consumers with a clear understanding of the sensitivity and access requirements for the data they are using.

F. User Interface and Experience

The data classification framework must provide a seamless and intuitive user experience to ensure effective adoption and utilization across the organization. This includes:

- Classification Tools:
    - Developing intuitive interfaces for manual classification and review processes is crucial for the success of the data classification framework.
    - These tools should allow users to easily assign classification labels, provide contextual information, and review the accuracy of automated classifications.
    - By creating a user-friendly experience, enterprises can encourage widespread adoption and ensure that the classification process is integrated into the daily workflows of data producers and consumers.
- Visualization Techniques:
    - To provide visibility into the classification status and trends across the organization, the framework should include dashboards and reporting capabilities.
    - These visualization tools can help data stewards, security teams, and compliance officers monitor the overall classification landscape, identify areas of concern, and track progress over time.
    - By presenting classification data in a clear and actionable manner, enterprises can make informed decisions about data governance, access management, and security initiatives.

Effective user interfaces and visualization techniques are essential for driving user adoption, maintaining data quality, and aligning the data classification framework with the organization's broader data management objectives. By investing in these user experience-focused components, enterprises can ensure that the classification system is seamlessly integrated into the daily operations of the business, enabling robust data governance and compliance.

G. Benefits and Metrics for Success

The implementation of a comprehensive data classification framework delivers a range of benefits for enterprises. Enhanced security is a key advantage, as the framework ensures that only authorized users can access sensitive data, reducing the risk of unauthorized access and potential data breaches. Improved compliance is another significant benefit, as the clear identification and categorization of regulated data enables organizations to more easily demonstrate adherence to various regulatory requirements, simplifying audits and reducing legal risks. The automation and standardization provided by the data classification framework can also lead to improved operational efficiency, significantly reducing the time and effort required for manual policy enforcement. Additionally, a unified data classification system enables enterprises to enforce consistent access control policies and data governance practices across their entire data landscape, regardless of the underlying data platforms or storage locations, ensuring governance consistency.

To measure the success of the data classification framework, enterprises can track several key metrics. Monitoring the reduction in unauthorized access incidents is a crucial indicator of the framework's effectiveness in protecting sensitive information. Measuring the percentage of correctly labeled data assets can help assess the reliability and consistency of the classification processes. The number of compliance audits passed without any violations is a direct indicator of the framework's ability to ensure regulatory compliance. Finally, tracking the time required to classify and protect newly ingested data can help evaluate the scalability and automation capabilities of the data classification solution, ensuring that it can keep pace with the growing volume of data. By closely monitoring these metrics, enterprises can continuously improve their data classification implementation, identify areas for enhancement, and demonstrate the tangible benefits of a robust data governance framework to stakeholders and decision-makers.

V. Policy-Based Access Control

A robust data classification framework must be closely integrated with the organization's access control mechanisms to ensure that data access is granted based on the assigned sensitivity levels. This policy-based approach to access management is crucial for protecting sensitive information and maintaining regulatory compliance.

A. Mapping Classifications to Access Policies

- Role-Based Access Control (RBAC)
    - In an RBAC model, data classifications are aligned with organizational roles and responsibilities. For example, customer service representatives may be granted access to "Confidential" customer data, while finance team members can access "Restricted" financial records.
    - By defining these role-based policies, enterprises can ensure that users can only access the data necessary for their job functions, reducing the risk of unauthorized exposure or misuse of sensitive information.
- Attribute-Based Access Control (ABAC)
    - ABAC policies go beyond just roles and incorporate other data attributes, such as classification levels, to make fine-grained access decisions.
    - For instance, a "Highly Sensitive" data classification may only allow access to a select group of executives or data stewards, even if they hold different job titles. The access is granted based on the specific attributes of the user and the data, rather than just their organizational role.

- This level of granularity in access control enables enterprises to implement more nuanced security measures and adapt to evolving data sensitivity requirements.

B. Dynamic Policy Enforcement

- Real-Time Policy Evaluation
  - The data classification framework should integrate with access control systems that can evaluate policies and make access decisions in real-time, based on the current data classification and user context.
  - For example, when a user attempts to access a dataset, the system can instantly check the classification level of the data and the user's attributes to determine if access should be granted or denied. This ensures that access control is enforced at the point of data access, rather than relying on static, pre-defined permissions.
- Policy Conflict Resolution
  - In complex enterprise environments, it's possible for multiple policies to apply to a single dataset or user. The data classification framework should include mechanisms to identify and resolve any potential conflicts between these policies.
  - For instance, if a user is assigned to both the "Finance" and "IT" roles, and the policies for these roles have different access permissions for a particular "Restricted" dataset, the framework should be able to determine the appropriate access level based on a predefined conflict resolution strategy (e.g., granting the most restrictive access).

C. Auditing and Monitoring

- Access Logs
  - Maintaining detailed access logs is crucial for compliance, security, and troubleshooting purposes. The data classification framework should integrate with logging systems to capture all data access attempts, including the user, data classification, access decision, and timestamp.
  - These comprehensive audit trails enable enterprises to investigate potential security incidents, demonstrate compliance with regulatory requirements, and identify areas for improvement in their access control policies.
- Anomaly Detection
  - By implementing anomaly detection systems, the data classification framework can identify unusual access patterns or potential policy violations, such as a user attempting to access data outside of their normal job responsibilities.
  - These anomaly detection capabilities allow enterprises to proactively monitor for security threats, investigate suspicious activities, and refine their access control policies to address any gaps or weaknesses.

By tightly integrating the data classification framework with the organization's access control policies, enterprises can ensure that sensitive information is only accessible to authorized users, in line with the assigned sensitivity levels. This holistic approach to policy-based access management is a critical component of the overall data governance and security strategy, enabling enterprises to protect their most valuable data assets while maintaining regulatory compliance.

VI. Data Quality and Classification

Maintaining high-quality data is crucial for the effectiveness of the data classification framework. This includes ensuring the completeness, accuracy, and consistency of the classified data. The framework addresses this through the following approaches:

A. Quality Metrics for Classified Data

- Completeness: Enterprises must ensure that all relevant data points are captured and classified. This means that the classification process covers the full scope of the organization's data assets, without any gaps or omissions.
- Accuracy: The framework must validate that the assigned classifications correctly reflect the nature and sensitivity of the data. This helps prevent misclassifications that could lead to security breaches or compliance issues.
- Consistency: Maintaining uniform classification across similar datasets is essential for enforcing consistent access control policies and data governance practices. Inconsistent classification can undermine the effectiveness of the overall framework.

B. Automated Quality Checks

- Rule-Based Validation: The framework includes automated checks and validations to verify the accuracy and consistency of the classification labels. This could involve predefined rules, such as checking for the presence of sensitive information in "Public" datasets or ensuring that all customer records are classified as "Confidential" or higher.
- Machine Learning for Quality Assurance: Advanced machine learning techniques can be leveraged to identify potential misclassifications or data quality issues. By training AI models on a corpus of correctly classified data, the framework can detect anomalies or outliers that may indicate classification errors.

C. Feedback Loops and Continuous Improvement

- User Feedback Mechanisms: The framework should provide a way for data consumers to flag potential classification errors or inconsistencies. This user feedback can be used to refine the classification rules and policies, ensuring that the framework remains accurate and up-to-date.
- Periodic Reviews: Enterprises should establish regular review processes to assess the effectiveness of the classification schema and access control policies. This allows for the identification of areas for improvement and the implementation of updates to address evolving data management requirements.

By implementing these data quality assurance measures, the data classification framework can maintain the integrity and reliability of the classified data, supporting robust data governance, access management, and compliance initiatives within the organization.

VII. Challenges and Future Directions

As enterprises continue to navigate the evolving landscape of data management, the data classification framework faces several key challenges and future directions:

A. Scaling Classification Processes:

As data volumes continue to grow, organizations must develop strategies to scale classification processes effectively. This may involve further advancements in automated classification techniques, such as leveraging more sophisticated machine learning models and natural language processing algorithms.

Enterprises will need to explore ways to streamline and optimize the classification workflows to keep pace with the rapid influx of new data.

B. Handling Unstructured Data:

While the current framework addresses the classification of structured data, improving techniques for handling unstructured data sources, such as documents, emails, and social media content, will be a crucial area of focus. Enterprises will need to invest in developing more advanced natural language processing and computer vision capabilities to accurately classify and categorize these diverse and complex data types.

C. Cross-Organizational Data Sharing:

As enterprises increasingly collaborate and share data across organizational boundaries, the data classification framework will need to address the challenge of maintaining classification integrity and consistent access control policies. Developing standardized protocols and frameworks for cross-organizational data sharing will be essential to ensure the protection of sensitive information and compliance with relevant regulations.

D. Ethical Considerations:

As the data classification framework becomes more reliant on automated decision-making, enterprises must address potential biases and fairness concerns within the classification algorithms and access control policies. Ensuring that the classification process and resulting access decisions are unbiased and equitable will be a critical area of focus, particularly as organizations navigate evolving privacy laws and societal expectations around data governance.

By addressing these challenges and exploring future directions, enterprises can continue to refine and enhance their data classification frameworks, ensuring that they remain effective, scalable, and aligned with the evolving needs of the modern, data-driven business landscape.

VIII. Conclusion

1. Data classification is a foundational element for enabling robust data governance, access management, and quality control in enterprise environments.
2. By implementing comprehensive data classification systems, organizations can:
    - Enhance their data security posture
    - Improve compliance efforts
    - Optimize resource allocation for data management initiatives
3. The integration of data classification with policy-based access control mechanisms provides a powerful framework for:
    - Protecting sensitive information
    - Ensuring appropriate access for authorized users
4. As data continues to grow in volume and importance, the role of data classification in enterprise data management will become increasingly critical.
5. Organizations that invest in developing and maintaining sophisticated data classification systems will be better positioned to:
    - Derive value from their data assets
    - Ensure regulatory compliance
    - Maintain a competitive edge in the data-driven business landscape
6. Future research directions for data classification include:

- Exploring advanced machine learning techniques for more accurate and efficient classification
- Developing standardized classification schemas for specific industries or data types
- Investigating novel approaches to handling the classification of rapidly evolving datasets in real-time environments

By adopting a comprehensive and scalable data classification framework, enterprises can secure sensitive data, improve operational efficiency, and ensure compliance with global regulations, positioning themselves for success in the increasingly data-driven business landscape.

REFERENCES

[1] R. Y. Wang and D. M. Strong, "Beyond Accuracy: What Data Quality Means to Data Consumers," J. Manage. Inf. Syst., vol. 12, no. 4, pp. 5–33, Mar. 1996. [Online]. Available: https://www.tandfonline.com

[2] C. Batini and M. Scannapieco, Data and Information Quality: Dimensions, Principles, and Techniques, Cham, Switzerland: Springer, 2016. [Online]. Available: https://link.springer.com

[3] J. Beech and F. Kriegel, "Metadata-Driven Approaches for Reusable Data Quality Metrics," in Proc. IEEE Int. Conf. Big Data (Big Data), Seattle, WA, USA, Dec. 2018, pp. 2581–2587. [Online]. Available: https://ieeexplore.ieee.org

[4] A. Al-Ruithe, R. Benkhelifa, and K. Hameed, "A Systematic Literature Review of Data Governance and Cloud Data Governance," Pers. Ubiquitous Comput., vol. 23, no. 5, pp. 839–859, Oct. 2019. [Online]. Available: https://link.springer.com

[5] F. Nargesian, E. Zhu, and R. J. Miller, "Data Lake Management: Challenges and Opportunities," Proc. VLDB Endowment, vol. 12, no. 12, pp. 1986–1989, Aug. 2019. [Online]. Available: https://vldb.org

[6] L. Cabral, T. Domingos, and E. Martins, "Data Lineage Management for Reproducible Science," in Proc. IEEE DSAA, Paris, France, Oct. 2015, pp. 1–10. [Online]. Available: https://ieeexplore.ieee.org

[7] T. Redman, "The Impact of Poor Data Quality on the Typical Enterprise," Commun. ACM, vol. 41, no. 2, pp. 79–82, Feb. 1998. [Online]. Available: https://dl.acm.org

[8] European Parliament and Council, "General Data Protection Regulation (GDPR)," Official Journal of the European Union, 2016. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2016/679/oj

[9] ISO/IEC, "ISO/IEC 27001:2013 - Information Security Management Systems," International Organization for Standardization, 2013. [Online]. Available: https://www.iso.org

[10] D. Loshin, Enterprise Knowledge Management: Data Governance and Compliance, Burlington, MA, USA: Morgan Kaufmann, 2010.