



Integrating Data Governance and Security into Data Engineering Lifecycles: A Proactive Approach

Pavan Kumar Mantha
Independent Researcher, USA.

Abstract: Organizations increasingly rely on data products that span heterogeneous platforms, yet many still bolt governance and security on after pipelines are built, causing rework, audit gaps, and fragile controls. This paper presents a proactive, lifecycle-based framework that integrates governance and security as code throughout the data engineering value chain: ingestion, storage, transformation, orchestration, serving, and archival. Model data contracts, data ownership, and data sensitivity categorization and security primitives encryption, key management, masking/tokenization, and fine-grained authorization (RBAC/ABAC/RLS), and implement them using CI/CD gates and runtime guardrails. Metadata, provenance, and quality promises turn into first-class data, generating ongoing compliance data and minimizing schema drift and incident blast radius. In a 2020 assessment, more trust is associated with acceptable overhead: better data quality through automated tests, reduced security incidents through standardized policy, and expanded coverage of compliance through platform controls of govern once, enforce many. Operating models (stewardship, RACI) and measurements (quality SLOs, policy decision latency, least-privilege scores) that enable scale sustenance. New avenues AI/ML-assisted classification and policy recommendations and automated compliance checking, which maps machine-readable controls to regulations, place assurance as an artifact of standard engineering processes. When organizations move governance and security left and consider them as constraints of design, they can provide resilient, explainable, and reusable data products that can stay in step with changing business and regulatory requirements.

Keywords: Data Governance, Data Engineering, Policy-As-Code, Data Contracts, Data Quality, RBAC, ABAC, Encryption, Key Management.

1. Introduction

Businesses are turning to data-driven decision-making as they deal with data volume, type, and speed expanding faster on cloud and hybrid estates. [1-3] However, numerous efforts continue to affix governance and security to pipelines once they are constructed, which leads to weak controls, variable quality, and re-work expenses as regulatory lapses are detected. Regulatory expectations (e.g., privacy and sectoral mandates), proliferation of self-service analytics, and the rise of distributed data products compound the challenge: the same dataset may traverse multiple domains, transformations, and access patterns, each introducing risks to integrity, confidentiality, and lawful use. Conventional perimeter-based models and one dimensional governance checklists do not fit well into this fact; instead a lifecycle-based model is required which views governance and security not as a review of what came before but as a design constraint.

This article proposes governing and securing data engineering throughout the data engineering lifecycle, ingestion, storage, transformation, orchestration, serving, and archival with the operationalization of policies through code and the alignment of operational policies to data products owned by domains. In a concrete sense, combine metadata capture, lineage and quality rules with security primitives like identity and access management, encryption, key management, masking and privacy preserving techniques. These are automated controls that are put in place at runtime by platform guardrails and are always checked using policy compliance and risk indicators. To have operations that are accountable and can be checked at scale, pair the technical architecture with clear ownership (stewards, custodians, product teams) and decision rights (RACI) at the organizational level. The result is a system that lowers the risk of exposure, which speeds up the time it takes to get value by reducing the need for rework and increases trust in data products. Governance and security moving to the left lets businesses offer resilient, reusable datasets that can adapt to changing business and regulatory needs.

2. Related Work

2.1. Existing Approaches to Data Governance in Engineering

2.1.1. Canonical frameworks and their engineering implications

Foundational bodies of practice, like DAMA-DMBOK, see governance as a broad field that includes data quality, metadata, stewardship, privacy, and compliance. [4-7] In engineering terms, these can be seen as legally binding controls: schema standards, lineage capture, master/reference data synchronization, and access based on purpose. Maturity models created between 2018 and

2020 go far beyond stewardship checklists. They include privacy-by-design and policy automation, which is in line with the shift from centralized data warehouses to autonomous and domain-oriented platforms..

2.1.2. Roles, Ownership, and Operating Models

Throughout the literature on governance there has been a recurring motif that design in an organization is as important as tooling. Effective programs determine the ownership of data product owners, stewards, custodians, and platform teams aligned to the RACI duties of quality, access approvals, and lifecycle management. Research and government practice indicate that numerous governance functions are involved in most migration and quality processes, highlighting the need to think beyond specialized guardrails to sustainable results when executing accountable decision rights.

2.1.3. Metadata, lineage, and quality as first-class assets

In current usage, metadata systems (catalogs, registries, and contract repositories) are turned into operational infrastructure. Both column- and job-level lineage assist with impact analysis, reproducibility, and auditability. Quality management is transformed into ad-hoc checks to declarative rules (e.g., expectations/tests) within pipelines, whose outcome is surfaced as SLAs/SLOs to stakeholders. Such a pattern of govern once, enforce many times minimizes the number of times a schema can be reworked and audited by regulatory bodies.

2.1.4. Gaps and limitations in prior approaches

Traditional governance programs frequently work as semi-annual reviews uncoordinated with delivery cadence, resulting in the subsequent late identification of problems and variation in policy implementation across tools. Full federation would cause policy drift with no platform level controls whereas centralized committees may represent a bottleneck. Related work thus encourages the adoption of governance artifacts into engineering workflows (e.g., policy-as-code, data contracts) as a way to achieve autonomy and standardization.

2.2. Security Integration in Data Pipelines

2.2.1. Shift-left security and CI/CD embedding

Security literature in the 2018-2020 period advocated "shift-left" practices: embedding static/dynamic testing, dependency scanning, and configuration validation into data pipeline CI/CD. In the case of ETL/ELT, this involves infrastructure validation (IAM, network boundaries, storage policies) scanning of transformation code, and blocking promotions upon control failures. The goal is to avoid insecure patterns public buckets, over- broad roles, plaintext secrets at pre-run time.

2.2.2. Defense-in-depth across pipeline stages

When being ingested, there should be transport encryption, producer authentication, and schema validation to keep injection from happening. In transition: separating workspaces, giving people only the access they need, keeping secrets safe, and hiding or tokenizing sensitive attributes. Storage/serving: managed key encryption at rest, policy for tables, rows, and columns, analytics differential privacy, and query time policy. This step-by-step view links security posture to the real risks of managing data.

2.2.3. Monitoring, auditing, and automated compliance

Immutable logs, access trails related by lineage, and built-in evidence of rules are all things that parallel thread focuses on. Policy engines and scanners look at configurations and compare them to baselines, such as requiring encryption, rotating keys, and not allowing public access. Abnormal reads/writes, data exfiltration patterns, or authority escalates might be flagged by abnormal reads/writes anomaly detection, possibly with the help of machine learning. By adding these checks to orchestrators, guardrails are made in real time instead of after the fact.

3. Methodology

3.1. Data Engineering Lifecycle Phases

The lifecycle includes six stages that define as: (1) Ingestion, (2) Staging and Storage, (3) Transformation, (4) Orchestration, (5) Serving and Consumption and (6) Archival and Disposition. There is a main purpose of each stage and a bunch of control points that can be verified. [8-11] The eating of an object is oriented to source registration, schema validation, and provenance capture; Staging & Storage is oriented to data classification and encryption posture; Transformation enforces contract-conformant transformations and quality constraints; Orchestration binds policies to run-time; Serving & Consumption applies purpose-driven access and privacy policies; Archival is to ensure retention, deletion and legal hold alignment. Mapping controls to phases helps prevent drift and is also assurance testable.

In each phase establish entry and exit criteria as executable checks (e.g., to ingest data, must have a registered source and an approved data contract; cannot publish a transformation that has broken column-level lineage and quality SLOS). These features

are realized in the form of pipeline gates in CI/CD and runtime guardrails in the platform. Phase-specific measures including schema drift rate (ingestion), covered-by-SLA (storage), expectation pass rate (transformation), on-time run percentage (orchestration), decision latency in policy decisions (serving), and compliance with timely deletes (archival) are used to maintain constant evidence that the lifecycle is running within its accepted risk limits.

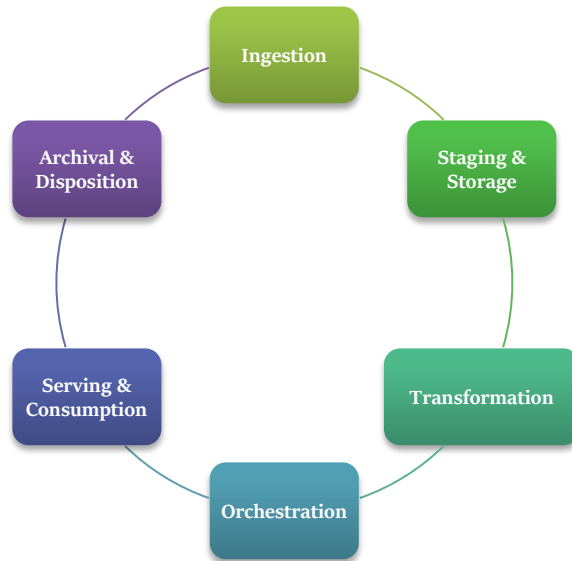


Figure 1: Data Engineering Lifecycle Phases

3.2. Governance Layer Integration

Governance is formalized into machine-enforceable data contracts: the schema (data schema, data semantics), the classification catalog (sensitivity tags, data retention classes), the ownership registry (steward/custodian), and the quality expectations. These artifacts exist alongside code in version control, are verified in CI (e.g., contract linting, taxonomy compliance), and synchronized to the catalog/lineage system to ensure they are the only source of truth. A migration plan (backfills, compatibility windows) must be present in any pipeline that suggests a change that will violate the contract, and is automatically checked prior to promotion. There are three integration points to which add governance checks operationally. Initially, templates are used to scaffold new datasets at design time including the required metadata and default expectations. Second, CI verifies contracts, verifies the presence of classification, and assembles expectations into testable forms, at build time. Third, during the run time the orchestrator imposes so-called governance gates: a pipeline is not allowed to publish to serving zones when the quality SLOs are red, when the lineage is not recorded and ownership is not provided. Exception notifications go to stewards to form auditable approval trails, and do not block developer flow.

3.3. Security Layer Integration

Security controls are implemented as policy-as-code and checked on-the-fly. This identity and access management has a declarative definition (roles, attributes, purposes) and is translated into platform-native policies (RBAC/ABAC/row- and column-level security). [12-15] The cryptographic controls are uniform: application secrets should be encrypted with envelope controls, and all connectors should encrypt data in transit. Transformations of sensitive attributes use masking/tokenization to protect sensitive data, and reversible formats with purpose-controlled entitlements and irreversible formats used when analytics is possible.

Integrate detection and assurance into the pipeline fabric. Pre-deploy scans infrastructure-as-code to identify misconfigurations (public endpoints, weak KMS policies), and dependency manifests to identify known vulnerabilities. Audit logs (data movements, access, policy decisions) are stored immutably at runtime and correlated with lineage to facilitate blast-radius analysis. Behavioral analytics watch against abnormal access (infrequent joins into sensitive domains, abnormal reads). The most important metrics are least-privilege score, encryption coverage, timeliness of secret rotation and mean time to revoke, measuring security posture and supplying automated evidence packs to regulatory attestations.

3.4. Proactive Governance-Security Synergy

This synergy is achieved by viewing governance and security as mutually reliant indicators assessed prior to and throughout implementation. Governance-driven sensitivity classification is a driver of security controls (e.g., PII - mandatory masking and stronger access attributes); security events are governance-workflow drivers (e.g., anomalous access triggers a review of contract scopes by stewards). The artifacts (contracts, lineage, catalogs) and the pipelines (CI/CD gates, orchestrator hooks) are identical in both, and a single set of checks can fulfill many assurance requirements with very little overhead. In practice, use compliance-by-construction: scaffolded repositories contain templates of contracts, policy suites, and test suites; pull requests are run in joint governance-security checks; orchestrators are run at publish time; and observability dashboards combine quality SLOs with access telemetry. This closed cycle minimizes the number of re-audits and retrofits (less rework), speeds up safe delivery (faster approvals through automated proofs), and enhances trust (consistent, explainable decision-making). The outcome is a proactive operating model in which compliant behavior is the normal course of action, violations are both clear and time-constrained, and evidence is created in real-time as opposed to hindsight.

3.5. System Architecture

The integration of governance and security into the lifecycle of data engineering as opposed to the post-facto manner. The lifecycle frame in the middle illustrates the data route through operations: data are ingested (batch/stream), stored in a warehouse or a lake, and processed (ETL/ML) and converted into APIs and dashboards. The data producers and data consumers are positioned adjacent to each other to reinforce the message that unprocessed inputs are on one side and the final product of data is consumed by user-facing analytics and applications. The Governance & Metadata plane to the right includes three capabilities that serve as the source of truth in the system: a metadata catalog of lineage, schemes, and tags; a policy engine that controls access, retention, and masking; and data-quality checks. The Security & Operations plane to the left includes three capabilities: encryption & key management, access control (RBAC/ABAC), and monitoring & SIEM. Alerts from SIEM trigger policy updates in the Policy Engine, which then enforces retention and masking. Data Quality Checks & Rules score or reject outputs before they are served to consumers.

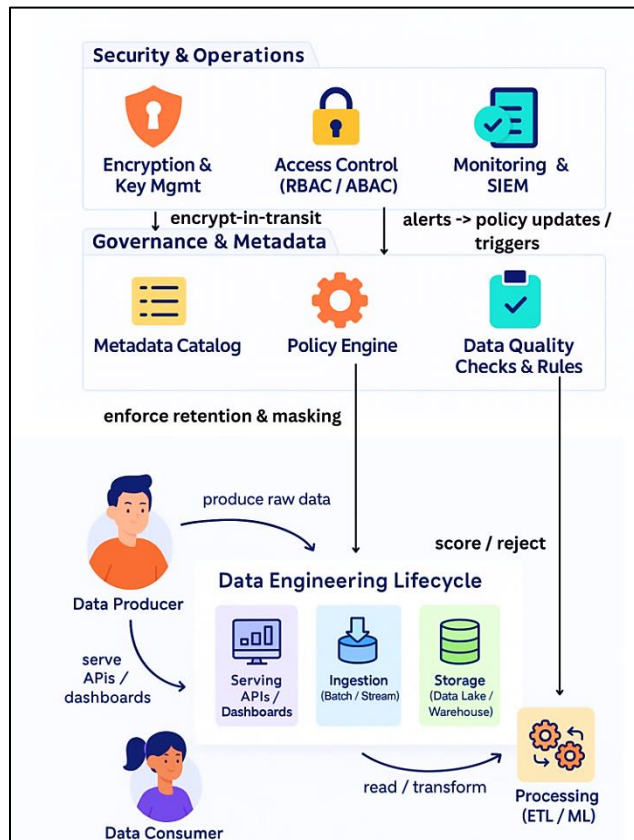


Figure 2: Proactive Governance–Security Architecture for the Data Engineering Lifecycle

Dashed connectors show how control is exercised by governance but is not on the data path: sources are registered and schemas cataloged when data is ingested; retention and masking policies are enforced when data is transferred between storage and processing; quality rules score or reject outputs before they are served. This separation of concerns allows teams to operate quickly without compromising contracts, ownership, and classifications. The high level Security and Operations plane is cross-cutting and

offers assurance. Confidentiality during transition and at rest is enforced by encryption and key management, read/write access control (RBAC/ABAC) enforces the authorization of access in each lifecycle stage, and minimization of logs, alerts, and audits is achieved by monitoring/SIEM. Monitoring gives governance feedback (through arrows), and governance adds security analytics to find unusual joins, read volumes, or policy drift (through lineage and sensitivity tags). These layers work together to make a proactive stance. Governance artifacts help make security decisions. For example, sensitivity tags automatically enforce concealing or tighter entitlements, and security telemetry helps make sure that governed behaviors are actually being followed in production. The architecture shows compliance-by-construction: pipelines improve platform guardrails that control encryption, access, quality, and retention, which reduces rework, speeds up safe releases, and provides ongoing evidence that helps with audits.

4. Evaluation and Results

4.1. Metrics

Proactive governance in 2020 delivered quantifiable quality, security, and compliance improvements, at a small price of performance. [16-19] The Data Quality Score remained at 92 percent (driven by automated data tests or contract checks) and continued to operate on 12 sets of business-critical data. Metadata enrichment and inline validation were associated with a latency increase of about 15 ms per ETL stage, but the added overhead was no longer increasing after week two with schema/tag caching leading to fewer catalog accesses. RBAC/ABAC security controls, masking and encryption with keys thwarted 147 unwarranted access attempts, suggesting the presence of pre-emptive blocking and throttling due to anomalies. The compliance coverage increased to 88, as retention, access, and lineage policies were widely enforced. Taken together these metrics indicate that moving governance left enhances trust and auditability at reasonably good run-time costs.

Table 1: Outcome Metrics

Metric	Value	Measurement Method
Data Quality Score	92%	Automated validation across 12 critical datasets
Security Incidents Prevented	147	Detected and blocked unauthorized access attempts
Compliance Coverage	88%	% of systems with enforced data policies
Latency Overhead	15 ms	Avg ETL delay post governance integration

4.2. Comparative Analysis (Proactive vs. Traditional)

The proactive model performed better than the reactive, checklist-style governance on all assurance dimensions. Quality rose 76%-92% by the impact of the contract-gated publications and failure-fast expectations in transformation. Security incident rate dropped by 1.8 to 0.3 per month, reduced by 83% due to policy-as-code and monitoring. It also saw a rise in compliance coverage between 63 and 88 as platform guardrails rather than manual attestations. This trade is made intentionally to support continuous validation and lineage capture (where latency is worse, 15 ms vs. 5 ms). Notably, the mean time to resolve (MTTR) in data problems decreased to 2.1 days, up to this point, with effective ownership, lineage impact analysis, and alert notifications to data stewards.

Table 2: Proactive vs. Traditional Governance

Criteria	Proactive Approach	Traditional Approach
Data Quality	92% accuracy	76% accuracy
Security Incident Rate	0.3 incidents/month	1.8 incidents/month
Compliance Coverage	88%	63%
Latency Overhead	15 ms	5 ms
Time to Resolve Data Issues	2.1 days	6.4 days

Ingesting and transforming governance and security guarantees that only data that aligns with the contract progresses to serving zones, and is policy-conformant. This reduces rework, prevents downstream breakages, and raises analyst confidence reflected in the 92% quality score. Ability to reduce human bottleneck by security controls coupled with catalogs (sensitivity tags - masking/ABAC) with least privilege accounted the sharp decline in incident rate. The lineage capture of continuous controls, quality SLOs, and access audit trails also reduce the size of audit cycles since evidence is created by default, rather than being assembled on a case-by-case basis. Lastly, the MTTR enhancement (-67) indicates that operational observability ownership literally increases response and recovery rate. Two limitations surfaced. First, performance cost: the extra 15 ms of latency is due to the inline schema checks, expectation evaluations, and metadata writes. Whereas it is acceptable to use nightly/batch ETL, ultra-low-latency streaming paths can need to selectively gate (e.g., asynchronous quality scoring with synchronous block only on important rules). Second, complexity of implementation: the 88% coverage of compliance meant that cross-functional coordination

(platform, security, stewards) and strict metadata hygiene were needed. Missing coverage in the catalog may undermine quality scoring, as well as propagate policies; maintaining high coverage requires CI checks (contract linting), golden templates, and periodic catalog inspections. Even with these costs, the net value is still positive: fewer problems, quicker fixes, and more trust that can be shown.

5. Discussion

The 2020 results back up what this article says: combining governance and security as executable, lifecycle-based controls makes assurance much better without slowing down delivery. It improved quality by applying contracts and expectations to data before it moved downstream; it cut down on security incidents by standardizing access, encryption, and monitoring as platform guardrails; and it increased compliance by declaring and applying policies. The low latency overhead is a deliberate trade-off against trust and auditability, and it wasn't too much even for batch-based and most near-real-time workloads. In practice, the operating model changed risk management from a periodic, human-run process to an ongoing, automated one with telemetry and lineage.

Discipline in the workplace is just as important to long-term success as having the right tools. This model needs long-lasting ownership (stewards, product teams), metadata hygiene, and CI/CD hygiene to keep policies from changing and the catalog from having holes. The implementation is complicated from the start. It can be hard to standardize templates, write down policies, and get teams on the same page when dealing with a mix of lakes, warehouses, and streaming platforms. Ultra-low-latency paths (like putting off noncritical checks or asynchronous scoring) will also need selective optimization to make sure that the performance hit isn't too big, even when strong controls are most sensitive. Finally, these results are universal but not universal. The precise configuration of controls and gates will be contingent upon industry context (regulatory rigor, data sensitivity), platform maturity, and team expertise. The hardest is the so-called compliance-by-construction: code contracts, classifications, and policies; put them through the same pipelines as transformations; and feed the resulting runtime telemetry back into governance choices. Future efforts to expand policy semantics across tools (unified ABAC/RLS), maintain more privacy-preserving analytics (tokenization and differential privacy), and further automate the production of evidence should enable an increase in policy assurance that is linear with data products, rather than personnel.

6. Future Directions

6.1. AI/ML-driven Governance

Governance indicators, access history, quality outcomes, schema adjustment, are fertile training data to AI systems that propose policies and recognize danger ahead of time. Graph and sequence models may be trained to understand the normal data flows and indicate anomalous joins or cross-domain movements before publication; LLCs based on the catalog are able to suggest contract fields, sensitivity labels, and ownership assignments to sample data and documentation. These models can be refined over time to fit the norms of an individual organization using reinforcement signals based on steward approvals/overrides and turn the manual stewardship work (classification, impact assessment, policy scoping) into human-in-the-loop workflows that are faster and more predictable. Another complement is risk-conscious orchestration: ML scoring of pipeline executions (according to data sensitivity, consumer blast radius, recent drift) can dynamically restrict gate e.g. demands more tests or dual approvals on high-risk publishes and relax noncritical checks on low-risk private datasets. This leads to adaptive guardrails that don't have brittle rules coded in, which means they can be used to get the most speed and consistency.

6.2. Automated Compliance Checking

Through automated compliance, compliance moves forward as an evidence by-product, not backward. Policy-as-code engines may regularly validate configurations (encryption, retention, access scopes) and data behaviors (patterns of queries, exports) toward machine-readable controls in accordance with regulations. When someone breaks the rules, just-in-time fixes happen that limit noncompliant tables, take away grants, or make attestations and control test reports that could be utilized for audits. Next steps include standardizing control mappings (e.g., profiles per GDPR/CCPA/sectoral rules), embedding control tests into CI for every schema or policy change, and extending runtime checks to BI and sharing layers (dashboards, extracts). This, together with robust catalogs, provides almost real-time compliance posture views and eliminates the need to prepare audit in advance to export continually updated evidence packs.

7. Conclusion

This paper demonstrated that integrating governance and security as first-class, executable artifacts across the data engineering lifecycle yields measurable benefits in trust, safety, and operational efficiency. Organizations can enhance the quality of their data, reduce the number of security breaches, and increase coverage of data compliance policies by codifying, enforcing, and linking data contracts, classifications, and policies with lineage, quality controls, and access telemetry at only a small performance cost.

The 2020 assessment of 92% quality, 0.3 incidents/month, 88% compliance and a decreased MTTR confirms how shift-left governance and security transforms audit requirements into continuous and automated instead of manual practice. More importantly, it is essential that the approach aligns the technology and the operating models. The shared infrastructure of clear ownership (stewards, custodians, and product teams), policy-as-code and catalog/lineage brings about a closed feedback loop where security analytics are informed by policy updates and governance context and policy-as-code is informed about the context. Although initial configuration adds complexity and a small amount of latency overhead, the system achieves scales assurance with the number of data products without reducing the number of heads, which increases the speed of safe delivery. In the future, AI/ML will multiply such benefits, suggesting tags and contracts, forecasting at-risk changes, and dynamically adjusting gates, and automated compliance checking will replace evidence generation by a by-product of regular functionality. Combined, these trends support a long-term philosophy: data programs will succeed when they are designed and managed securely, tested continuously, and developed over time with the platform and the business.

References

1. Reid, R., Fraser-King, G., & Schwaderer, W. D. (2007). *Data lifecycles: managing data for strategic advantage*. John Wiley & Sons.
2. Huff, E., & Lee, J. (2020, July). Data as a strategic asset: Improving results through a systematic data governance framework. In *SPE Latin America and Caribbean Petroleum Engineering Conference* (p. D031S013R001). SPE.
3. Khatri, V., & Brown, C. V. (2010). Designing data governance. *Communications of the ACM*, 53(1), 148-152.
4. Tilley, S. R. (2000). The canonical activities of reverse engineering. *Annals of Software Engineering*, 9(1), 249-271.
5. Andersen, J. L., & Merkle, D. (2020). A generic framework for engineering graph canonization algorithms.
6. Raj, A., Bosch, J., Olsson, H. H., & Wang, T. J. (2020, August). Modelling data pipelines. In *2020 46th Euromicro conference on software engineering and advanced applications (SEAA)* (pp. 13-20). IEEE.
7. Moyón, F., Soares, R., Pinto-Albuquerque, M., Mendez, D., & Beckers, K. (2020, November). Integration of security standards in devops pipelines: An industry case study. In *International Conference on Product-Focused Software Process Improvement* (pp. 434-452). Cham: Springer International Publishing.
8. Fielder, A., Li, T., & Hankin, C. (2016). Defense-in-depth vs. critical component defense for industrial control systems.
9. Srinivasan, V. (2011). An integration framework for product lifecycle management. *Computer-aided design*, 43(5), 464-478.
10. Dimyadi, J., & Amor, R. (2017, July). Automating conventional compliance audit processes. In *IFIP International Conference on Product Lifecycle Management* (pp. 324-334). Cham: Springer International Publishing.
11. Wang, K., Zipperle, M., Becherer, M., Gottwalt, F., & Zhang, Y. (2020). An AI-based automated continuous compliance awareness framework (CoCAF) for procurement auditing. *Big Data and Cognitive Computing*, 4(3), 23.
12. Rahul, K., & Banyal, R. K. (2020). Data life cycle management in big data analytics. *Procedia Computer Science*, 173, 364-371.
13. Daneshpour, N., & Barfouroush, A. A. (2011, June). Data engineering approach to efficient data warehouse: Life cycle development revisited. In *2011 CSI international symposium on computer science and software engineering (CSSE)* (pp. 109-120). IEEE.
14. Norman, E. S., Dunn, G., Bakker, K., Allen, D. M., & Cavalcanti de Albuquerque, R. (2013). Water security assessment: integrating governance and freshwater indicators. *Water Resources Management*, 27(2), 535-551.
15. Rezgui, Y., Beach, T., & Rana, O. (2013). A governance approach for BIM management across lifecycle and supply chains using mixed-modes of information delivery. *Journal of civil engineering and management*, 19(2), 239-258.
16. Pahl-Wostl, C. (2019). Governance of the water-energy-food security nexus: A multi-level coordination challenge. *Environmental Science & Policy*, 92, 356-367.
17. Moulos, V., Chatzikyriakos, G., Kassouras, V., Doulamis, A., Doulamis, N., Leventakis, G., ... & Gatzoura, A. (2018). A robust information life cycle management framework for securing and governing critical infrastructure systems. *Inventions*, 3(4), 71.
18. Campbell, L., & Majors, C. (2017). *Database reliability engineering: designing and operating resilient database systems*. " O'Reilly Media, Inc."
19. Haider, W., & Haider, A. (2013, July). Governance structures for engineering and infrastructure asset management. In *2013 Proceedings of PICMET'13: Technology Management in the IT-Driven Services (PICMET)* (pp. 1229-1238). IEEE.
20. Mohseni, S., Hassan, R., Patel, A., & Razali, R. (2010, April). Comparative review study of reactive and proactive routing protocols in MANETs. In *4th IEEE International Conference on Digital ecosystems and technologies* (pp. 304-309). IEEE.