

From Dashboards to Decisions: A Microsoft Fabric and Copilot-Enabled Enterprise Analytics Governance Framework for Saudi Vision 2030

CHOUDHRY BILAL MAZHAR HUSSAIN

Digital Analytics & AI Lead, Arabian Agricultural Services Company (ARASCO), Riyadh, Saudi Arabia
ORCID: [0009-0005-7625-7296](https://orcid.org/0009-0005-7625-7296)

Abstract Enterprises are moving from descriptive dashboards toward governed, AI-assisted decision environments. Microsoft Fabric unifies data engineering, warehousing, real-time analytics and business intelligence on OneLake, while Microsoft Copilot introduces natural-language decision support across the analytics lifecycle. Yet unified data and generative AI do not automatically produce trustworthy decisions. This paper develops and validates by case application a design-science framework that connects Fabric and Copilot with Saudi Arabia's National Data Management Office (NDMO), National Data Index (NDI), Personal Data Protection Law (PDPL), and Vision 2030 priorities. The framework contributes to a five-layer architecture, a governance-to-decision maturity model, a Copilot-in-the-loop workflow, a capability-to-control mapping, and decision-ready KPIs. Practice-based validation is provided through an ARASCO enterprise analytics case involving existing Power BI, staging SQL database, and governance-dashboard contexts. The central argument is that enterprise value is created not by dashboards or generative AI alone, but by governing the path from question to validated decision through lineage, sensitivity labeling, role-based access, human decision rights and audit evidence.

Index Terms- Data Governance; Decision Support; Enterprise Analytics; Microsoft Copilot; Microsoft Fabric; Responsible AI

I. INTRODUCTION

Saudi Arabia's Vision 2030 positions data and artificial intelligence as national enablers of economic diversification, government efficiency and quality of life. SDAIA describes itself as the Kingdom's

competent authority for data and AI, and its National Data Index frames data-management maturity, compliance and operational excellence as measurable public-sector capabilities (SDAIA 2025a, 2025b; Saudi Press Agency 2023). Within this context, enterprise analytics is no longer a back-office reporting function. It is part of the information infrastructure through which organizations govern performance, risk, compliance, and innovation.

Yet enterprises often remain data-rich and decision-poor. Dashboards, reports, data extracts, and spreadsheet workarounds accumulate faster than trusted decision pathways. The problem is not the absence of visualization, but the absence of a governed route from business question to data asset, from data asset to verified insight, and from insight to accountable action. This distinction is especially important when generative AI is embedded directly in the analytics workflow. A Copilot-generated narrative or DAX expression may accelerate analysis, but it is not a decision unless it is grounded in governed by semantic models, validated by a responsible actor and recorded as evidence.

Part of the problem is architecture. Enterprise analytics has traditionally depended on separately managed ingestion tools, data lakes, warehouses, engineering platforms, reporting layers and governance services. Microsoft Fabric responds to this fragmentation by integrating data engineering, warehousing, real-time intelligence, data science and business intelligence into a single software-as-a-service platform built around OneLake, a unified logical data lake using open Delta Parquet storage and a tenant, workspace and domain hierarchy for federated ownership (Microsoft 2026a, 2026b). Copilot then adds natural-language

exploration and generative assistance over the analytics lifecycle (Microsoft 2026c, 2026d).

This paper asks: How can a Microsoft Fabric and Copilot-enabled enterprise analytics environment be governed so that it moves organizations from dashboards to trustworthy, auditable and decision-ready outcomes aligned with Saudi Vision 2030? The contribution is a design-science framework supported by practice-based validation in an ARASCO case setting. The paper contributes: (1) a five-layer Fabric-Copilot reference architecture; (2) a governance-to-decision maturity model; (3) a Copilot-in-the-loop workflow; (4) a capability-to-control matrix aligned to Saudi data governance; (5) a decision-ready KPI dictionary; and (6) a case-based validation showing how the artefacts map to a real enterprise analytics modernization context.

II. RELATED WORK AND CONCEPTUAL BACKGROUND

2.1 Knowledge systems, analytics and decision support

Knowledge and advanced information systems research has long addressed decision support, knowledge processing, data sharing, data warehousing, interoperability, visualization, and intelligent information retrieval (Shim et al. 2002). These themes position KAIS as a natural outlet for a paper that treats dashboards, semantic models and Copilot as components of a decision-oriented knowledge system rather than isolated business-intelligence tools. Decision-support research also shows that analytics changes organizational routines when analytical capabilities are embedded in information-processing processes, not when analytical outputs remain disconnected from decisions (Shollo and Galliers 2016; Wieder and Ossimitz 2015; LaValle et al. 2011).

Business intelligence and analytics literature have consistently emphasized that value is realized through managerial use, strategy alignment, and organizational maturity. Davenport and Harris (2007) argue that analytics becomes strategic when it contributes to competitive differentiation. Popovic et al. (2012) show that business-intelligence system maturity contributes to information quality and use. Langer (2025)

synthesizes data and analytics maturity models and differentiates organization-oriented, technology-oriented, and data-oriented maturity dimensions. Cardoso and Su (2022) show how maturity-model artefacts can be developed and evaluated using design science. The proposed framework extends this stream by focusing on AI-augmented analytics in a Fabric and Copilot environment.

2.2 Data governance, data products and semantic control

Data-governance literature has moved beyond policy documentation toward operational control. Weber et al. (2009) propose a contingency approach that links governance arrangements to organizational context. Khatri and Brown (2010) identify decision domains in data governance, including data principles, quality, metadata, access, and lifecycle. Otto (2011) develops requirements for mature data governance, while Alhassan et al. (2016) provides a systematic review of data-governance activities. More recent reviews show that implementation barriers remain substantial even when governance frameworks exist (Acuña et al. 2026; Guerreiro et al. 2025; Systematic Analysis of Data Governance Frameworks 2025).

Data-governance literature has moved beyond policy documentation toward operational control. Weber et al. (2009) propose a contingency approach that links governance arrangements to organizational context. Khatri and Brown (2010) identify decision domains in data governance, including data principles, quality, metadata, access, and lifecycle. Otto (2011) develops requirements for mature data governance, while Alhassan et al. (2016) provides a systematic review of data-governance activities. More recent reviews show that implementation barriers remain substantial even when governance frameworks exist (Acuña et al. 2026; Guerreiro et al. 2025; Acev et al. 2025).

2.3 Responsible AI and AI governance for analytics

Responsible AI governance literature emphasizes that AI systems require accountability, transparency, human oversight, privacy protection, and lifecycle monitoring. The NIST AI Risk Management Framework structures AI risk management around governance, mapping, measurement and management (Tabassi 2023). Reviews of AI governance stress that many frameworks articulate principles but provide less

guidance on operational mechanisms and organizational accountability (Batool et al. 2023; Tangi et al. 2024; Cheong 2024). For enterprise analytics, this gap is practical: a Copilot-assisted answer needs to be verified, access-controlled, and explainable before it can influence a business decision.

The present paper contributes to an operational bridge between data governance and AI governance. It does not evaluate a machine-learning model. Instead, it governs the knowledge path through which natural language, semantic models, generated measures, human validation, and audit evidence interact. This is particularly relevant in Saudi Arabia, where NDMO, NDI and PDPL requirements create a policy environment in which data quality, classification, compliance and operational excellence are explicit governance objectives (SDAIA 2025b; NDMO 2025; Kingdom of Saudi Arabia 2023).

III. MATERIALS AND METHODS

This study follows a design-science research approach. Design science creates artefacts that address relevant organizational problems through constructs, models, methods, or instantiations (Hevner et al. 2004; Peffers et al. 2007). The artefacts were developed through a structured narrative synthesis and validated by application to an enterprise case setting. This is not a statistical generalization study; it is an artefact-development and practice-validation paper.

The evidence base consisted of three classes. First, Microsoft documentation described platform capabilities for Fabric, OneLake, Direct Lake, Copilot, and Purview (Microsoft 2026a, 2026b, 2026c, 2026d, 2026e, 2026f). Second, Saudi policy artefacts described NDMO, NDI, PDPL and Vision 2030 data-governance priorities (SDAIA 2025a, 2025b; NDMO 2025; Kingdom of Saudi Arabia 2023). Third, peer-reviewed literature grounded data governance, maturity models, decision support, design science and AI governance (Acuña et al. 2026; Guerreiro et al. 2025; Langer 2025; Tabassi 2023).

The case validation draws on ARASCO's project materials and correspondence related to Microsoft

Fabric, Power BI, data warehousing, and governance dashboards. Only generalized facts are reported, and no commercially sensitive server identifiers, contractual values, personal data, credentials or confidential operational measures are disclosed. The case is used for validation by application: the proposed artefacts are mapped against the observed enterprise analytics current state, proposed target architecture, governance requirements and implementation boundary conditions. During manuscript preparation, Microsoft 365 Copilot was used only for language editing and formatting assistance; it was not used to generate research findings, data or analytical conclusions. The author reviewed all generated texts and takes responsibility for the final content.

IV. FRAMEWORK RESULTS

The synthesis produced five artefacts: a five-layer reference architecture, a governance-to-decision maturity model, a Copilot-in-the-loop decision workflow, a decision-ready KPI dictionary, and a responsible-AI/PDPL control matrix. The artefacts are intended to be evaluated together because platform architecture without governance controls does not produce trustworthy decisions, and governance controls without decision metrics do not demonstrate value.

4.1 Five-layer reference architecture

The reference architecture in Figure 1 organizes a Fabric and Copilot estate into five capability layers bound by a vertical governance-and-trust spine. OneLake forms the unified foundation. Fabric workloads provide ingestion, lakehouse, warehouse, real-time and data-science capabilities. Governed semantic models and Direct Lake provide certified business logic and low-latency access. Copilot adds natural-language analytics, DAX assistance, and narrative summaries. Role-based decision surfaces deliver insight to executives, operations teams, and self-service users. Across all layers, Purview, sensitivity labels, RBAC, RLS, lineage, NDMO evidence and responsible-AI controls constrain data use and support accountability.




Figure 1. Microsoft Fabric and Copilot-enabled enterprise analytics governance reference architecture.

Table 1. Platform capability-to-governance-control mapping.

Capability / layer	Function	Governance control
OneLake storage	Single logical data foundation	Tenant policies, domain ownership, classification
Workspaces and domains	Federated ownership and administration	RBAC, workspace roles, capacity governance
Semantic models	Certified business logic and measures	Endorsement, certification, measure descriptions
Direct Lake	Low-latency analytics without unnecessary copies	Lineage and refresh governance
Copilot natural-language layer	Question-to-insight interaction	AI-ready content scoping and RLS

DAX/narrative generation	Generated calculations and summaries	Human validation and verified answers
Fabric data agents	Conversational access to governed data	Decision rights and audit logging

Saudi alignment	Evidence artefact
NDMO cataloguing and classification	Data catalogue and classification register
NDMO governance roles and accountability	Access-review and ownership register
NDMO data quality; NDI maturity	Model certification and measure catalogue
NDMO lineage and interoperability	Lineage graph and refresh log
PDPL data minimization and access control	Prompt, response and access log
Responsible AI; NDMO quality	Verification and validation record
PDPL accountability; NDI evidence	Immutable decision audit trail

4.2 Governance-to-decision maturity model

Figure 2 describes progression from fragmented reporting to auditable decision support. The model deliberately places Copilot-enabled augmentation after the governed stage. This sequencing matters: introducing generative AI over uncertified content may increase the speed of answer generation, but it does not improve trust. Table 2 summarizes the maturity stages.




Figure 2. Governance-to-decision maturity model for Fabric and Copilot adoption.

Table 2. Governance-to-decision maturity stages.

Stage	Characteristics	Governance posture
Siloed	Fragmented tools and duplicated data	Ad hoc and inconsistent
Managed	Consolidated on Fabric/OneLake	Basic roles and workspaces
Governed	Certified models, labels and lineage	Enforced classification and access
Augmented	Copilot in the decision loop	Human validation and decision rights
Auditable	Evidence-backed decision support	Continuous assurance and audit

Copilot role	Indicative maturity measure
None or ungoverned	Low catalogue and quality scores
Limited and unscoped	Rising catalogue coverage
Restricted to AI-ready content	Improved quality and lineage coverage

Insight, DAX and narrative assistance	Verification and adoption of metrics tracked
Routine governed use	High maturity across governance dimensions

4.3 Copilot-in-the-loop decision workflow

Figure 3 operationalizes the path from question to action. A decision-maker asks a natural-language question, which is resolved against a certified semantic model rather than an unmanaged data store. Copilot generates an insight, DAX expression or narrative explanation. The output then passes through human validation and decision-rights checks before action is taken. The action, input, model version and rationale are logged, and feedback is used to refine models, measures and AI-ready content scope.




Figure 3. Copilot-in-the-loop decision workflow with human validation and audit feedback.

4.4 Decision-ready KPIs and responsible-AI controls

Table 3 replaces report-centric indicators with decision-centric indicators. A dashboard estate that produces many reports but does not reduce decision latency or improve verified insight adoption is not mature. Table 4 translates responsible-AI and privacy risks into controls that preserve explainability, privacy and audibility before Copilot outputs influence operational or executive decisions.

Table 3. Decision-ready KPI dictionary.

KPI	Definition / formula	Data source
Decision latency	Time from question to validated decision	Workflow and audit logs
Insight adoption rate	Validated insights acted upon / total generated	Decision workflow logs

Copilot verification rate	Verified Copilot outputs / outputs used in decisions	Copilot usage logs
Semantic model certification coverage	Certified models / total models in use	Fabric/Purview catalogue
Sensitivity-label coverage	Labelled critical assets / total critical assets	Purview
Data quality score	Composite of completeness, accuracy and consistency	Quality platform
Self-service adoption	Active governed self-service users / eligible users	Fabric usage logs
Measure reuse	Reused certified measures / total measures	Semantic model catalogue

Frequency	Owner	Evidence
Weekly	Analytics lead	Decision timeline export
Monthly	Domain owner	Adoption register
Weekly	AI governance owner	Verification log
Monthly	Data governance office	Certification record
Monthly	Data protection officer	Label coverage report
Monthly	Data stewards	Quality scorecard
Monthly	BI lead	Usage report
Quarterly	Data governance office	Reuse register

Table 4. Responsible-AI and PDPL control matrix for Copilot-assisted decisions.

Risk principle	Fabric/Copilot mechanism	Governance action	Evidence
Hallucination or incorrect output	Verified answers and certified semantic models	Require human validation before action	Verification record

Sensitive-data exposure	Sensitivity labels, RLS and scoped AI-ready content	Restrict Copilot to permitted content	Access and label logs
Explainability	Measure descriptions and narrative rationale	Require analyst-readable justification	Explanation log and model card
Human oversight	Decision-rights gate	Enforce human-in-the-loop for high-impact actions	Approval record
Auditability	Purview lineage and audit logs	Record input, model version, output and action	Immutable audit trail
PDPL lawful basis	Data minimization and access control	Apply lawful-basis and privacy checks	Data-protection impact note

V. CASE VALIDATION: ARASCO ENTERPRISE ANALYTICS MODERNIZATION

5.1 Case rationale and current-state observations

The ARASCO case was selected because it represents a typical regulated enterprise analytics environment in which operational dashboards, staging databases, governance reporting and emerging AI expectations coexist. Internal project materials describe a current state including Microsoft 365 licensing with Power BI Pro, multiple GCP and on-premises staging database servers, Power BI Report Server, Power BI Gateway and no existing Fabric subscription. The proposed target state in those materials includes procuring Fabric capacity, consolidating staging data into OneLake, retiring or reducing legacy ETL and reporting infrastructure where validated, and enabling AI, Copilot and governed analytics.

This current state reflects the problem formulation of the paper. The immediate issue is not a lack of dashboards; it is a fragmented route from enterprise source systems to trusted decision outputs. Separate staging databases and report servers can support dashboards, but they complicate lineage, certification, semantic reuse and AI-ready governance. The case therefore provides a practical setting to test whether the framework captures actual implementation constraints.

5.2 Mapping the framework to the ARASCO target architecture

The ARASCO Fabric proof-of-concept materials describe a target flow from SAP S/4HANA to Fabric Data Factory, then to OneLake medallion layers and Power BI Direct Lake dashboards. They also identify governance capabilities such as Microsoft Purview lineage, sensitivity labeling, policy enforcement, Entra ID access integration and SDAIA/NDMO alignment. These elements map directly to the five-layer

architecture: SAP and other enterprise systems form the source layer, Fabric Data Factory and pipelines form the ingestion layer, OneLake medallion architecture forms the governed data layer, semantic models and Direct Lake form the analytics layer, and executive or operational dashboards form role-based decision surfaces.

The case also validates the need for decision-ready KPIs. Internal materials identify executive dashboards, procurement reporting and cost dashboards as candidate value-sprint reports. These are not merely technical dashboards; they are decision surfaces for governance, procurement, infrastructure and executive performance. The framework therefore adds decision latency, verification rate, semantic certification and evidence coverage as measures of value because report availability alone would not prove improved decision support.




Figure 4. ARASCO practice-based validation mapping from current analytics environment to proposed Fabric-Copilot governance architecture.

5.3 Evidence from completed governance dashboard implementation

A completed ARASCO governance-dashboard closure document provides additional validation. It reports that a FreshService governance system was configured with service catalogues, workspaces, automated workflows and request items, and that a Power BI governance dashboard was deployed using API-based integration and automated refresh. The same document identifies lessons learned that align closely with the proposed framework: define the governance data model and workflow structure upfront, align early with subject-matter experts and system owners, plan API integration and data mapping during design, and conduct early training for faster adoption.

This completed governance-dashboard case does not validate Microsoft Fabric specifically. It validates the broader mechanism proposed by the paper: governance analytics succeeds when workflows, ownership, API integration, dashboard interpretation and training are designed together. It also shows why the framework treats decision surfaces, governance processes and adoption controls as one socio-technical system.

5.4 Validation findings and boundary conditions

Table 5. Case-based validation of the framework against ARASCO implementation evidence.

Framework artefact	ARASCO case evidence	Validation interpretation	Boundary condition
Reference architecture	Proposed SAP S/4HANA → Fabric Data Factory → OneLake medallion → Power BI Direct Lake flow	Architecture matches a plausible enterprise target state	POC and procurement remain required before full deployment
Governance spine	Purview lineage, sensitivity labels, Entra ID access and SDAIA/NDMO alignment identified in internal materials	Governance controls are not optional add-ons but architectural requirements	Exact control implementation depends on tenant configuration
Maturity model	Current state includes fragmented staging/reporting assets; target state consolidates on Fabric	Siloed-to-governed progression fits observed context	Maturity level is qualitative, not statistically measured
Copilot workflow	Fabric/Copilot is positioned as an AI-enabled analytics capability requiring business and systems-team approval	Human validation and decision rights are required before automation	Copilot usage must be limited to permitted and governed content
Decision-ready KPIs	Candidate dashboards include executive, procurement and infrastructure/cost reporting	Decision KPIs are needed to show business value beyond dashboard count	Baseline values were not collected in this study
Implementation controls	FreshService governance dashboard lesson: define data models, workflows, SME alignment and API mapping early	Lessons support workflow-first governance design	Evidence comes from one organization and one completed governance-dashboard project

The validation confirms that the framework is practically applicable to an enterprise analytics modernization context. It does not prove performance improvement. No before-and-after quantitative measurement was conducted for Fabric deployment because Fabric was still positioned as a procurement and POC decision in the reviewed materials. The case

therefore provides validation by fitness and implementation plausibility, not causal evidence. This distinction is important for scientific transparency and for avoiding overstatement.

VI. DISCUSSION

The framework reframes Fabric and Copilot adoption as a move from data platform to decision platform. A data platform organizes storage, pipelines and reports. A decision platform connects data, governance, AI assistance, human validation and action. This distinction matters because dashboards alone create visibility without accountability. The case validation supports the same view: ARASCO already had dashboards and reporting infrastructure, but the strategic question was how to consolidate data, improve governance and enable AI-ready decision support.

Governance is a binding constraint. Data unification reduces fragmentation, but it also centralizes risk. Copilot accelerates inquiry, but it can amplify errors when content is not certified, labelled and contextually understood. The Saudi context strengthens the governance requirement because NDMO, NDI and PDPL turn data governance from a best-practice concept into a measurable and legal responsibility. The ARASCO case shows how these obligations appear in practical architecture decisions: Purview lineage, sensitivity labels, Entra ID integration and SDAIA/NDMO alignment are treated as architectural controls, not documentation afterthoughts.

The framework also suggests a shift in performance measurement. Traditional analytics teams track report counts, refresh success, model accuracy or user logins. These metrics are useful but incomplete. Decision latency, verification rate and insight adoption measure whether analytics improves decisions. For executive adoption, what matters is not how many dashboards exist but whether the organization can make trusted decisions faster and with better evidence.

The framework is Saudi-specific in its control mapping but internationally transferable in its architecture. Organizations outside Saudi Arabia could adopt the same five-layer architecture and decision workflow while replacing NDMO, NDI and PDPL with their own governance and privacy obligations. This preserves regional relevance without limiting conceptual generalizability.

VII. CONCLUSION

This paper developed a Microsoft Fabric and Copilot-enabled enterprise analytics governance framework aligned with Saudi Vision 2030, SDAIA, NDMO, NDI and PDPL priorities. The framework contributes a reference architecture, maturity model, Copilot-in-the-loop workflow, control mapping, decision-ready KPI dictionary and responsible-AI/PDPL control matrix. It also adds practice-based validation using an ARASCO enterprise analytics modernization case.

The main conclusion is that enterprise value is not created by dashboards or generative AI alone. Value arises when the path from business question to validated decision is governed through lineage, semantic certification, sensitivity labels, access control, human decision rights and audit evidence. The case validation shows that these artefacts fit a real enterprise context involving existing Power BI, staging databases, governance dashboards and a proposed Fabric target architecture.

The study has limitations. The validation is based on one organization, and quantitative before-and-after results were not available. Future work should conduct expert evaluation with data-governance leaders, Fabric architects, AI governance specialists and Saudi compliance practitioners; then pilot the framework in one business domain and measure changes in decision latency, semantic model certification coverage, Copilot verification rate and insight adoption. Pending such validation, the framework provides an implementation-ready reference for regulated enterprises seeking to move from dashboards to auditable, AI-assisted decision support.

DECLARATIONS

Funding: This research received no external funding.

Ethical Approval: Not applicable.

Informed Consent: Not applicable.

Data Availability: No new public dataset was created or analyzed. The case validation uses generalized observations from internal enterprise project materials. Confidential details, system identifiers, credentials and financial values are not disclosed.

Competing Interests: The author declares no competing interests.

Author contributions Conceptualization, methodology, formal analysis, investigation, writing—original draft preparation, writing—review and editing, and visualization: Choudhry Bilal Mazhar. Use of generative AI During preparation of this manuscript, Microsoft 365 Copilot was used for language editing and formatting assistance. The author reviewed and edited the output and takes full responsibility for the final content.

REFERENCES

- [1] Acev D, Biyani S, Rieder F, Aldenhoff TT, Blazevic M, Riehle DM, Wimmer MA (2025) Systematic analysis of data governance frameworks and their relevance to data trusts. *Management Review Quarterly*. [Crossref](#)
- [2] Alhassan I, Sammon D, Daly M (2016) Data governance activities: a comparison between scientific and practice-oriented literature. *Journal of Enterprise Information Management* 29(2):300–316. [Crossref](#)
- [3] Acuña I, Cravero A, Zorrilla M (2026) From compliance to value: data governance implementation challenges to turn data into a strategic asset—a systematic literature review. *Applied Sciences* 16:5289. [Crossref](#)
- [4] Batool A, Zowghi D, Bano M (2023) Responsible AI governance: a systematic literature review. arXiv:2401.10896. <https://doi.org/10.48550/arXiv.2401.10896>
- [5] Cardoso E, Su X (2022) Designing a business intelligence and analytics maturity model for higher education: a design science approach. *Applied Sciences* 12:4625. [Crossref](#)
- [6] Cheong BC (2024) Transparency and accountability in AI systems: safeguarding wellbeing in the age of algorithmic decision-making. *Frontiers in Human Dynamics* 6:1421273. [Crossref](#)
- [7] Davenport TH, Harris JG (2007) *Competing on analytics: the new science of winning*. Harvard Business School Press, Boston.
- [8] Dehghani Z (2022) *Data mesh: delivering data-driven value at scale*. O'Reilly Media, Sebastopol.
- [9] Dolhopolov A, Castelltort A, Laurent A (2024) Implementing federated governance in data mesh architecture. *Future Internet* 16:115. [Crossref](#)
- [10] Guerreiro L, Martins J, Bernardo MR, Mamede H, Branco F (2025) Impact of a master data management framework to trigger data governance maturity: a systematic literature review. *IEEE Access* 13:190644–190667. [Crossref](#)
- [11] Hevner AR, March ST, Park J, Ram S (2004) Design science in information systems research. *MIS Quarterly* 28:75–106. [Crossref](#)
- [12] Hupperz M, Ghosh A, Nelson ML (2024) Data products, data mesh, and data fabric. *Business & Information Systems Engineering* 66:529–536.
- [13] Khatri V, Brown CV (2010) Designing data governance. *Communications of the ACM* 53(1):148–152. [Crossref](#)
- [14] Kingdom of Saudi Arabia (2023) Personal Data Protection Law (PDPL). Riyadh, Saudi Arabia.
- [15] LaValle S, Lesser E, Shockley R, Hopkins MS, Kruschwitz N (2011) Big data, analytics and the path from insights to value. *MIT Sloan Management Review* 52(2):21–32.
- [16] Langer B (2025) Understanding data & analytics maturity: a systematic review of maturity model composition. *Management Review Quarterly*.
- [17] Microsoft (2026a) OneLake, the unified data lake—Microsoft Fabric. Microsoft Learn. <https://learn.microsoft.com/en-us/fabric/onelake/>
- [18] Microsoft (2026b) Fabric architecture for a unified data platform—Cloud Adoption Framework. Microsoft Learn.
- [19] Microsoft (2026c) Copilot for Power BI overview. Microsoft Learn.
- [20] Microsoft (2026d) Copilot in Power BI integration overview. Microsoft Learn.
- [21] Microsoft (2026e) Direct Lake overview—Microsoft Fabric. Microsoft Learn. <https://learn.microsoft.com/en-us/fabric/fundamentals/direct-lake-overview>
- [22] Microsoft (2026f) Microsoft Purview data governance documentation. Microsoft Learn.
- [23] NDMO (2025) National data management and personal data protection standards. Saudi Data and AI Authority, Riyadh.
- [24] Otto B (2011) Organizing data governance: findings from the telecommunications industry and consequences for large service providers. *Communications of the Association for Information Systems* 29:45–66. [Crossref](#)

- [25] Peffers K, Tuunanen T, Rothenberger MA, Chatterjee S (2007) A design science research methodology for information systems research. *Journal of Management Information Systems* 24:45–77. [Crossref](#)
- [26] Popovic A, Hackney R, Coelho PS, Jaklic J (2012) Towards business intelligence systems success: effects of maturity and culture on analytical decision making. *Decision Support Systems* 54(1):729–739. [Crossref](#)
- [27] Saudi Press Agency (2023) SDAIA, NTP launch first National Data Index, and open data, data governance platforms. Saudi Press Agency, Riyadh.
- [28] SDAIA (2025a) Saudi Data and AI Authority. Saudi Data and AI Authority, Riyadh. <https://sdaia.gov.sa/en/default.aspx>
- [29] SDAIA (2025b) National Data Index (NDI). Saudi Data and AI Authority, Riyadh. <https://sdaia.gov.sa/en/Research/Pages/NationalDataIndex.aspx>
- [30] Shim JP, Warkentin M, Courtney JF, Power DJ, Sharda R, Carlsson C (2002) Past, present, and future of decision support technology. *Decision Support Systems* 33(2):111–126. [Crossref](#)
- [31] Shollo A, Galliers RD (2016) Towards an understanding of the role of business intelligence systems in organisational knowing. *Information Systems Journal* 26(4):339–367. [Crossref](#)
- [32] Tabassi E (2023) Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. [Crossref](#)
- [33] Tangi L, Janssen M, Benedetti M, Noci G (2024) AI governance in a complex and rapidly changing regulatory landscape. *Humanities and Social Sciences Communications* 11:1130. [Crossref](#)
- [34] Weber K, Otto B, Österle H (2009) One size does not fit all—a contingency approach to data governance. *Journal of Data and Information Quality* 1(1):1–27. [Crossref](#)
- [35] Wieder B, Ossimitz ML (2015) The impact of business intelligence on the quality of decision making—a mediation model. *Procedia Computer Science* 64:1163–1171. [Crossref](#)