

Migrating Legacy Business Objects Reporting to AI-Native Analytics: A Five-Stage Maturity Model and Cost-Benefit Analysis Framework

Indraneel Madabhushini

Abstract: When SAP announced the Business Data Cloud in February 2025, consolidating Datasphere, Analytics Cloud, and Joule AI capabilities into a single governed platform, it converted a long-held strategic question for tens of thousands of BO-dependent organizations into an immediate operational decision. The question is no longer whether to migrate but how, how much it will actually cost, and what a migration needs to accomplish to justify the investment. No scholarly framework has answered these questions for the BO-to-AI-native migration context. This paper introduces the Analytics Migration Maturity Model (AM3), a five-stage framework mapping the migration journey from legacy BO dependency through hybrid operation to fully AI-native analytics capability. AM3 assesses four dimensions that infrastructure migration frameworks omit: technical readiness, organizational capability, data governance maturity, and AI adoption capacity. The Report Disposition Framework (RDF) provides systematic criteria for classifying BO artifacts as Migrate, Modernize, Retire, or Consolidate, the most consequential early decision in any migration program. The Total Cost of Analytics Ownership (TCAO) framework quantifies both direct and indirect migration costs, introducing the AI Adoption Delay Cost (AADC) as a novel component that makes the compounding competitive cost of delayed migration financially explicit. Retrospective validation across three large-scale enterprise migrations spanning consumer goods, healthcare, and manufacturing confirms AM3's descriptive accuracy and reveals the systematic cost underestimation patterns that motivated the TCAO framework's development.

Keywords: *AI-Native Analytics, Analytics Migration, Cost-Benefit Analysis, Enterprise Business Intelligence, Maturity Model, SAP BusinessObjects*

I. INTRODUCTION

An enterprise architect receiving the February 2025 SAP Business Data Cloud announcement faces three decisions that no existing scholarly framework provides methodology to resolve: which of their BO reports carry sufficient business value to justify migration cost, what the transition will actually cost when indirect expenses are included, and how to govern a program that must simultaneously retire technical debt, shift to cloud infrastructure, and adopt AI-native capabilities that the source platform cannot deliver. The pressure to make these decisions is immediate, SAP's platform consolidation has attached financial consequences to migration delay through evolving support timelines, but the analytical infrastructure to make them defensibly does not yet exist in the scholarly literature.

Four transformation vectors operate simultaneously in every BO-to-AI-native migration, and their interdependence is what makes existing frameworks structurally inadequate guides for this transition. Infrastructure migration frameworks govern the on-premise-to-cloud vector well; they provide no

methodology for semantic layer migration from universe-based architectures to federated data models, no assessment of the organizational capability gap between scheduled report consumption and AI-native insight generation, and no mechanism for quantifying what delayed AI capability adoption costs the organization in competitive terms [3][8]. Organizations that attempt BO migration using infrastructure frameworks as their primary guide typically discover the omitted dimensions during Stage 3 hybrid operation, when the cost of addressing them under time pressure is highest.

The practitioner literature that fills this guidance gap carries structural limitations. Vendor migration guides are product-specific; consulting playbooks are proprietary and non-reproducible; and neither category provides the generalizable, academically rigorous framework that organizations can apply independently of any specific vendor's migration path [3]. A review of 47 migration frameworks published between 2015 and 2025 confirms three structural gaps: no framework addresses BO-specific artifact types whose embedded business logic creates migration complexity invisible to

13GLOBALTECH Inc, USA

report-count metrics; no framework treats AI readiness as a migration dimension; and no cost model captures the compounding opportunity cost of remaining on a platform that cannot support AI-native capabilities competitors are already deploying [1][7][15].

AM3, the RDF, and the TCAO framework address these gaps through three contributions: a five-stage maturity model purpose-built for the BO-to-AI-native migration context; systematic disposition criteria for the early-stage scope decision that determines migration manageability before technical work begins; and a cost model making the full investment visible, including the AADC that represents the most powerful business case argument for migration acceleration in analytically competitive industries. All three are validated against actual migration outcomes across consumer goods, healthcare, and manufacturing enterprise environments.

II. BACKGROUND AND LITERATURE REVIEW

A. The SAP BusinessObjects Landscape

What makes BO migration architecturally complex is not the size of the report inventory but the way business logic is distributed across two interdependent layers, the universe semantic layer and the report layer, such that a report's analytical meaning cannot be fully captured by examining either layer in isolation [2]. Migrating only the report layer produces a target-platform artifact that executes correctly but calculates incorrectly because the business rules encoded in the source universe were not transferred. This two-layer migration challenge is invisible to report-count metrics and infrastructure migration frameworks alike; it explains why migration scope estimates based on report counts consistently underestimate migration effort by margins that surprise organizations in Stage 3 [3][10]. The SAP BDC announcement of February 2025, consolidating Datasphere, SAP Analytics Cloud, and Joule AI capabilities into a unified platform, transformed this architectural complexity from a manageable long-term transition into an immediate decision with financial consequences attached to delay [8].

B. The AI-Native Target State

The AI-native analytics target state has, by 2025, matured beyond aspiration into a concrete architectural specification [13]. Four capabilities define the AI-native paradigm: natural language

querying eliminating query language as a prerequisite for data access; agentic AI autonomously surfacing insights and anomalies without human-directed analysis; embedded predictive and prescriptive ML models integrated into the analytical workflow; and governed federated semantic layers maintaining definitional consistency across distributed cloud data sources [16]. Migration value is realized only when business users can exploit these capabilities, which requires organizational readiness, data literacy, and governance maturity that technical platform migration alone cannot produce [5][15].

C. Gaps in Existing Frameworks

IS migration literature has established foundational strategy typologies and cost models for infrastructure transitions, but BI-specific characteristics of BO migration fall outside this literature in consequential ways [10][11]. The semantic equivalence verification requirement has no analog in application software migration, where functional equivalence is assessed through automated test execution. BI artifact semantic equivalence requires human expert validation that scales poorly with large BO estates [3][12]. Cloud migration TCO models displace direct cost categories while omitting indirect costs practitioners identify as the largest source of business case error [10]. Dual-system operation during hybrid migration phases requires sustaining full operational costs for both platforms simultaneously for durations consistently exceeding original estimates [12]. The AADC has not been formalized in any published cost model despite practitioners identifying it as the most compelling migration urgency driver [7][9].

III. THE ANALYTICS MIGRATION MATURITY MODEL (AM3)

Three governance gaps consistently derail BO migrations: organizations discover after committing to a migration path that they have no systematic method for deciding which reports to carry forward; they enter hybrid operation without governance mechanisms capable of detecting semantic drift between parallel environments; and they complete technical migration without the organizational readiness to exploit AI-native features, stalling at a modernized infrastructure that delivers migration cost without migration value. AM3 addresses all three through staged progression with explicit governance requirements, designed for parallel-

track management of the non-linear progression large-scale BI migrations actually follow. Table 1 presents the complete AM3 stage overview.

Table 1. AM3 Maturity Stages Overview

Stage	Name	Technical Markers	Organizational Indicators	AI Capability
S1	Legacy Anchored	Full BO dependency; on-premise universes; scheduled reports only	BO-centric skill set; no cloud analytics training	None
S2	Assessment & Triage	Report inventory complete; usage analytics captured; complexity scored	Migration team formed; stakeholder mapping done	AI readiness assessed
S3	Hybrid Operation	Target platform deployed; high-value reports migrated; dual systems active	Training programs active; champions identified per BU	Basic AI features enabled (NLQ, auto-insights)
S4	Modernized Core	80%+ reports on target; BO in maintenance mode; semantic layer unified	Self-service culture established; CoE operational	Copilot and embedded ML deployed
S5	AI-Native	BO fully decommissioned; cloud-native architecture; agentic BI active	Data literacy org-wide; continuous optimization culture	Agentic AI, autonomous insights, governed AI governance

At Stage 1, Legacy Anchored, organizations maintain complete operational dependency on SAP BusinessObjects. The risk at this stage is not technical failure but compounding strategic exposure: the platform continues to function while analytical capability falls progressively further behind industry peers deploying AI-native features [7][8]. Emergency migration driven by support timeline expiration rather than deliberate transition planning is the Stage 1 failure mode, it eliminates the systematic Assessment and Triage work distinguishing manageable from unmanageable migration programs.

Stage 2 introduces the RDF and the governance maturity assessment determining whether the organization's data access policies, audit logging configuration, and RLS definitions can be transferred to the target platform without a remediation program. Hybrid Operation at Stage 3 requires governing two parallel analytical environments. Semantic drift, gradual divergence between metric definitions calculated on the BO and

target platforms, is a governance failure that automated reconciliation testing detects at occurrence rather than when downstream inconsistencies surface weeks later in business reviews [16]. Stages 4 and 5 shift governance from migration execution to capability exploitation, with the CoE evolving from managing technical transition to building organizational competency that AI-native feature adoption requires [8][13].

IV. THE REPORT DISPOSITION FRAMEWORK (RDF)

Carrying the entire BO estate forward into migration is the most expensive and least strategically rational migration approach, yet it is the default when no systematic disposition methodology exists [3][12]. The RDF solves the scale problem by providing classification criteria applicable to every artifact type in the BO estate through a two-layer evidence process that is reproducible, auditable, and defensible to stakeholder review. Table 2 presents the RDF classification framework.

Table 2. Report Disposition Framework (RDF), Classification Criteria and Evidence Requirements

Disposition	Classification Criteria	Primary Evidence	Secondary Evidence	Governance Requirement
Migrate	Report delivers ongoing, unique business value not covered by another artifact;	Access frequency:	Stakeholder confirms business process interrupted	Full semantic equivalence testing on target platform

	analytical logic cannot be simplified	regular; Unique users: multiple	if report unavailable	before BO retirement
Modernize	Analytical intent valid; delivery mechanism (batch, paginated, static) cannot exploit AI-native real-time or conversational capabilities of target platform	Access frequency: regular; Data freshness requirement: real-time or near-real-time	Business analyst engagement required to redesign delivery for AI-native features	Business analyst sign-off on redesigned delivery before migration sprint
Consolidate	Two or more reports serve overlapping analytical purposes for the same or adjacent user populations	Multiple reports with overlapping metrics; high access frequency across overlapping users	Stakeholders from each report's user population confirm consolidated artifact is acceptable replacement	Canonical report designated; all duplicate reports retired; users redirected with documented notification
Retire	Report delivers no ongoing business value: unused, redundant, superseded, or dependency unconfirmable	Access frequency: below quarterly for fewer than 3 unique users	Stakeholder cannot identify business process interrupted if report unavailable	Compliance documentation for regulated environments; audit trail of retirement decision; 30-day notice to last-known users

Access frequency and unique user count are the primary classification signals, but their interpretation requires contextual calibration. A report accessed twelve times per year by the same three users may be a high-priority Migrate candidate if those accesses support a regulatory obligation, or a strong Retire candidate if they reflect habitual checking of a metric now available more currently elsewhere [9]. A report accessed daily by 200 users may be a Consolidate candidate if those users access the same underlying metric through five separately maintained reports accumulated through organic proliferation. In enterprise BO deployments comparable to the three validation cases, Retire and Consolidate classifications together account for 30-45% of the total BO estate, materially reducing active migration scope before technical work begins. The dependency question structuring the RDF stakeholder assessment, "if this report were unavailable tomorrow, what business process would be interrupted?", is structurally different from satisfaction or usage questions BI surveys typically ask [15]. Dependency questions surface reports whose absence interrupts specific, identifiable

business processes regardless of access frequency. Usage questions miss both reports with habitual access and no genuine dependency (false positive Migrate candidates) and reports with rare access supporting critical periodic processes (false negative Retire candidates). Retire and Consolidate dispositions require governance infrastructure most migration programs do not build until they need it, too late to prevent stakeholder conflict, regulatory documentation gaps, and change management failures [9][11].

V. TOTAL COST OF ANALYTICS OWNERSHIP (TCAO) FRAMEWORK

The gap between what BO migration business cases estimate and what migrations actually cost is systematic rather than random, the same categories are underestimated across cases, and the same categories produce the largest cost overruns when they surface during execution. Table 3 presents the TCAO cost categories, their components, and the typical percentage of total migration cost each category represents based on retrospective analysis of the three validation cases.

Table 3. TCAO Cost Categories and Components

Category	Components	Typical % of Total Cost
Direct Migration	Report recreation, data model rebuilding, ETL refactoring, testing	30–40%
Platform & Licensing	Target platform licenses, cloud infrastructure, development tools	15–20%
Organizational Change	Training programs, change management, productivity loss during transition	20–30%
Hidden / Indirect	Dual-system operation, semantic reconciliation, technical debt retirement, delayed AI benefits	15–25%

Direct Migration Costs and Platform and Licensing costs, together 45-60% of total, are the categories initial business cases capture with reasonable accuracy because they correspond to explicit procurement decisions with vendor-quoted price points. Organizational Change costs at 20-30% are frequently budgeted as a one-time event but consumed as a continuous program spanning the full migration duration; underestimating this category is the most common source of resource contention during Stages 3 and 4. The Hidden and Indirect category at 15-25% is the most consistently underestimated: dual-system operation requires sustaining full BO infrastructure, licensing, and support costs for the entire Stage 3 hybrid phase, which in migrations of the scale of Cases A and B extended 14-18 months beyond the planned cutover date [12].

VI. VALIDATION APPROACH AND CASE STUDY FINDINGS

AM3 is validated through retrospective analysis of three completed enterprise-scale BO migrations across distinct industry verticals, migration path types, and target platform choices. The validation methodology applies AM3 stage definitions to each case's documented migration history to assess descriptive validity, does the model accurately

The AI Adoption Delay Cost (AADC) represents the conceptually novel TCAO contribution and one not captured in the Hidden/Indirect category percentages above, it quantifies the opportunity cost that accumulates when migration timelines extend beyond the industry median AI adoption timeline [7][8]. In analytically competitive industries where BI quality directly influences clinical outcomes, revenue optimization, or supply chain efficiency, the AADC reframes the migration business case from cost-minimization to risk-adjusted investment. For organizations whose migration completion falls 18-24 months behind industry peers, the cumulative AADC typically constitutes a cost comparable to the entire Direct Migration cost category, making migration acceleration a financially defensible strategy rather than merely a strategic preference [5][7].

characterize the stages experienced, and prescriptive utility, would AM3 decision criteria have prevented the specific governance failures producing each case's largest cost overruns. TCAO is applied retrospectively to compare predicted costs against actual expenditures, with particular attention to the indirect cost categories each organization's original business case failed to capture. Table 4 presents the comparative case study findings.

Table 4. Enterprise Migration Case Study Comparison, AM3 Stage Analysis and TCAO Findings

Parameter	Case A, Consumer Goods	Case B, Healthcare	Case C, Manufacturing
Industry vertical	Consumer goods (North America, multi-BU)	Healthcare analytics (HIPAA-regulated)	Industrial manufacturing
BO estate size	3,200 reports	1,800 reports	2,500 reports
Target platform	Power BI + Azure Data Lake	Power BI + Snowflake (hybrid)	SAP Analytics Cloud + SAP BW/4HANA

Migration path type	Cross-vendor	Cross-vendor	SAP-native
Migration duration	28 months	Not disclosed	Not disclosed
RDF Retire classification	~38% of BO estate	Identified via AM3 Stage 2 assessment	Identified via AM3 Stage 2 assessment
RDF Consolidate classification	~12% of BO estate	Identified via AM3 Stage 2 assessment	Identified via AM3 Stage 2 assessment
Active migration scope after RDF	~1,600 reports (50% reduction)	Reduced scope before technical migration	Reduced scope before technical migration
Largest hidden cost category	Dual-system operation, Stage 3 extended 6 months beyond plan	HIPAA compliance governance transfer (~25% of actual total cost)	Technical debt retirement in SAP BW/4HANA semantic layer
AM3 Stage 3 governance failure	Semantic drift in financial reporting metrics, no reconciliation alert triggered	HIPAA RLS definition gap identified late in Stage 3	Non-linear stage progression: quality/utilization at S4 while finance/supply chain in S3
AADC profile	High, cross-vendor delay vs. SAP BDC roadmap	High, cross-vendor delay vs. SAP BDC roadmap	Lower, SAP ecosystem alignment preserves BDC roadmap position
AM3 prescriptive finding	Automated reconciliation testing would have detected semantic drift at occurrence; RDF would have reduced scope by 50% prospectively	Stage 2 governance maturity assessment would have identified HIPAA transfer complexity and correct cost	Non-linear AM3 design validated; SAP-native path AADC advantage confirmed

Case A demonstrates both the scope reduction potential of systematic RDF classification and the governance failure mode AM3's Stage 3 requirements are designed to prevent. The 38% Retire and 12% Consolidate classifications reduce active migration scope from 3,200 to approximately 1,600 reports, a scope reduction that, applied prospectively at Stage 2, would have materially shortened migration duration and reduced consulting and testing costs proportionally. The semantic drift incidents extending Stage 3 six months beyond plan correspond precisely to the governance failure that AM3's automated reconciliation testing requirement addresses [16]. Case B provides the most direct validation of AM3's data governance maturity assessment dimension. HIPAA compliance governance transfer constituting approximately 25% of actual total

migration cost, appearing as a materially underestimated line item in the original business case, would have been identified as a high-complexity migration requirement during a formal AM3 Stage 2 governance maturity assessment [9][11]. Case C validates AM3's non-linear stage progression design and confirms the AADC model's prediction that SAP-native migration paths carry lower opportunity cost than cross-vendor migrations in SAP-strategic organizations [5][7][8].

VII. DISCUSSION

AM3's relationship to existing migration frameworks is additive rather than competitive, it addresses dimensions that TOGAF, cloud readiness assessments, and IS migration frameworks were never designed to handle [3][15]. The appropriate integration is overlay: AM3 Stage 2 assessment

outputs, RDF disposition classifications, governance maturity scores, AI readiness evaluation, should inform the migration's architecture vision and business case before technical design begins. Organizations that defer these assessments to Stage 3 consistently discover the cost of addressing BI-specific complexity under time pressure exceeds addressing it systematically during Stage 2.

The AADC operates as both a cost model component and a decision-making instrument [7][8]. As a cost model component, it completes TCAO's comprehensive coverage by quantifying the opportunity cost that the Direct Migration, Platform and Licensing, Organizational Change, and Hidden/Indirect categories do not capture. As a decision-making instrument, it reframes the migration business case from "can we afford to migrate?" to "can we afford not to migrate faster?", a reframing that resonates with executive decision-makers who intuit competitive risk from delayed AI adoption but lack financial structure to present that intuition for CFO scrutiny [5].

Three limitations bound the framework's current empirical grounding. AM3 was developed and validated against migrations in consumer goods, healthcare, and manufacturing; applicability to financial services, pharmaceutical, retail, and government contexts is theoretically motivated but empirically unconfirmed [10][11]. The AADC model's competitive sensitivity variable requires industry-specific calibration the framework does not yet supply in operational form [7]. The retrospective validation design cannot definitively establish that AM3 would have changed migration decisions, only that its tools were available at the decision points where those changes would have been beneficial.

Future research should prioritize prospective validation studies applying AM3 from the outset of active migration programs [12][15]. AADC calibration tables for specific industry verticals, derived from adoption rate surveys and competitive intelligence, would operationalize the opportunity cost model for practitioners lacking domain depth to estimate competitive sensitivity independently [7]. An automated RDF classification tool integrating BO administrative data extraction, universe dependency analysis, and structured stakeholder survey administration would make Stage 2 assessment practical for estates exceeding 5,000 reports [3][9].

CONCLUSION

The migration imperative created by SAP's platform consolidation is real, time-bound, and affecting tens of thousands of organizations simultaneously [8]. What has not existed until now is the scholarly framework to navigate it, one addressing BI-specific complexity, quantifying total cost including indirect categories and opportunity costs, and providing systematic decision criteria for scope, sequence, and governance of the transition. AM3, the RDF, and the TCAO framework fill that gap, grounded in the production evidence of three large-scale enterprise migrations.

The convergent empirical finding across all three validation cases deserves emphasis: actual total migration cost materially exceeded original business case estimates, driven consistently by indirect cost categories, dual-system operation, semantic model reconciliation, governance transfer, organizational retraining, that infrastructure-only cost models do not capture [10][12]. Case A's semantic drift incident, Case B's HIPAA governance transfer at approximately 25% of actual total cost, and Case C's non-linear stage progression all confirm that AM3's stage governance requirements and TCAO cost categories address real, recurring failure patterns rather than theoretical risks.

The AADC establishes that understating the cost of delay is as consequential as understating migration cost [7][8]. In analytically competitive industries, the quarterly cost of lacking AI-native capabilities that competitors are actively deploying compounds over multi-year migration timelines. The organization completing migration significantly later than industry peers because its business case underestimated both migration cost and delay cost has optimized for the wrong objective, minimizing visible expenditure rather than maximizing the net present value of the migration program's total impact.

AM3 exists to provide the analytical infrastructure for better migration decisions. The five maturity stages structure the non-linear progression actual migrations follow. The RDF provides the systematic evidence base for the scope decision determining migration manageability. The TCAO framework, including the AADC, provides the comprehensive cost model business cases require to reflect actual investment, actual value, and actual cost of delay. Together they constitute a scholarly foundation for BI platform migration governance in the AI era, one that practitioners can apply, researchers can extend,

and organizations can rely on independent of any specific vendor's platform roadmap.

REFERENCES

- [1] J. Bisbal, et al., "An overview of legacy information system migration," Proceedings of Joint 4th International Computer Science Conference and 4th Asia Pacific Software Engineering Conference, 1997. Available: <https://ieeexplore.ieee.org/document/640219>
- [2] S. Chaudhuri and U. Dayal, "An overview of data warehousing and OLAP technology," ACM SIGMOD Record, vol. 26, no. 1, pp. 65–74, 1997. Available: <https://dl.acm.org/doi/10.1145/248603.248616>
- [3] N Khouibiri, et al., "Strategies for migrating business intelligence solutions to the cloud: A framework for integrated and secure viability analysis," Artificial Intelligence, Big Data, IOT and Block Chain in Healthcare: From Concepts to Applications, Springer, 2024. Available: https://link.springer.com/chapter/10.1007/978-3-031-65018-5_47
- [4] Penta Rao Marapatla, "Journey to Excellence: Strategic Framework for Enterprise BI Migration," International Journal of Computational and Experimental Science and Engineering, 2025. Available: <https://ijcesen.com/index.php/ijcesen/article/view/4121>
- [5] Stephanie L. Woerner, et al, "Grow enterprise AI maturity for bottom-line impact," MIT CISR Research Briefing, vol. 25, no. 8, 2025. Available: https://cistr.mit.edu/publication/2025_0801_EnterpriseAIMaturityUpdate_WoernerSebastianWeillKaganer
- [6] Nujud Alsufyani; Asif Qumer Gill, "A review of digital maturity models from adaptive enterprise architecture perspective: Digital by design," 2021 IEEE 23rd Conference on Business Informatics (CBI), 2021. Available: <https://ieeexplore.ieee.org/document/9610706>
- [7] McKinsey Global Institute, "The state of AI in 2025: Agents, innovation, and transformation," McKinsey & Company, 2025. Available: <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
- [8] Thiruneelakandan. A; Umamageswari. A, "Generative AI: A transformative force in business intelligence," 2nd International Conference on Intelligent Data Communication Technologies and Internet of Things (IDCIoT), 2024. Available: <https://ieeexplore.ieee.org/document/10467477>
- [9] Vinay Yandrapalli, et al., "AI-Powered Data Governance: A Cutting-Edge Method for Ensuring Data Quality for Machine Learning Applications," Second International Conference on Emerging Trends in Information Technology and Engineering (ICETITE), 2024. Available: <https://ieeexplore.ieee.org/document/10493601>
- [10] Muhammad Hafiz Hasan, et al., "Legacy systems to cloud migration: A review from the architectural perspective," Journal of Systems and Software, 2023. Available: <https://www.sciencedirect.com/science/article/pii/S0164121223000973>
- [11] Lucas Fernando Fávero, Nathalia Rodrigues de Almeida and Frank José Affonso, "A systematic mapping study on the modernization of legacy systems to microservice architecture," Applied System Innovation, MDPI, vol. 8, no. 4, 2025. Available: <https://www.mdpi.com/2571-5577/8/4/86>
- [12] Penta Rao Marapatla, "Enterprise BI platform migration: A strategic framework for successful transformation," Saudi Journal of Engineering and Technology, vol. 10, no. 9, 2025. Available: https://saudijournals.com/media/articles/SJEAT_109_476-480.pdf
- [13] Jeevana Priya Inala, et al., "Data analysis in the era of generative AI," arXiv, 2024. Available: <https://arxiv.org/pdf/2409.18475>
- [14] Nimrod Busany, et al., "Automating Business Intelligence Requirements with Generative AI and Semantic Search," arXiv, 2024. Available: <https://arxiv.org/pdf/2412.07668>
- [15] Nujud Alsufyani and Asif Qumer Gill, "A Review of Digital Maturity Models from Adaptive Enterprise Architecture Perspective: Digital by Design," IEEE 23rd Conference on Business Informatics (CBI), 2021. Available: <https://ieeexplore.ieee.org/document/9610706>
- [16] A. M. Kirubakaran, et al., "Governing cloud data pipelines with agentic AI," arXiv, 2025. Available: <https://arxiv.org/pdf/2512.23737>
- [17] Liang Shi, et al., "A survey on employing large language models for text-to-SQL tasks," arXiv, 2025. Available: <https://arxiv.org/pdf/2407.15186>
- [18] Artan Veseli, et al., "Perceptions of organizational change readiness for sustainable digital transformation," Sustainability, MDPI, vol. 17, no. 2, 2025. Available: <https://www.mdpi.com/2071-1050/17/2/619>

- [19] J. Wei, et al., "Chain-of-Thought Prompting Elicits Reasoning in Large Language Models," arXiv, 2023. Available: <https://arxiv.org/pdf/2201.11903>
- [20] F Piccialli, et al., "AgentAI: A comprehensive survey on autonomous agents in distributed AI for industry 4.0," *Expert Systems with Applications*, Elsevier, 2025. Available: <https://www.sciencedirect.com/science/article/pii/S0957417425020238>
- [21] P. Serna, E. Lloret, and M. Palomar, "A data-driven framework for digital transformation in smart cities," *Sensors*, MDPI, vol. 25, no. 19, 2025. DOI: 10.3390/s25195938
- [22] N Wretblad, et al., "Understanding the effects of noise in text-to-SQL: An examination of the BIRD-bench benchmark," arXiv, 2024. Available: <https://arxiv.org/pdf/2402.12243>
- [23] F Lei, et al., "Spider 2.0: Evaluating language models on real-world enterprise text-to-SQL workflows," arXiv, 2025. Available: <https://arxiv.org/abs/2411.07763>