

Pillar 5: AI Use Cases - The Orchestration and Approval System for AI-Ready Data Governance

Gregory David Spehar
GiDanc AI LLC
myVibecoder.us
Version 1.1 Copyright ©2025

Background: AI use case orchestration serves as the approval and provisioning system in AI-ready data governance, addressing project failures in mid-sized enterprises. **Problem:** Unvetted AI initiatives contribute to \$10-14 million annual losses with 74% struggling to scale. **Method:** This paper presents pattern-based orchestration with federated approval processes. **Contributions:** Demonstrates 40-60% reduction in AI project failures and 20-30% efficiency gains through structured intake. **Implications:** Enables ethical AI deployment with regulatory compliance for 500-5,000 employee organizations. **Type:** Position Paper.

Keywords: AI use cases, orchestration, model governance, AI readiness, mid-sized enterprises

Introduction

AI use case orchestration serves as a critical component of AI-ready governance frameworks, managing requests, approvals, and data provisioning to ensure projects are compliant and successful. This pillar addresses AI-specific challenges by assessing data readiness and tailoring preparation patterns (e.g., LLM/RAG for natural language, predictive ML for forecasting), incorporating regulations like BCBS 239 for assessments, GDPR for privacy, and DORA for resilience (Deloitte Team, 2025; N-iX, 2024).

Production implementations in enterprise data platforms demonstrate significant operational improvements, with organizations achieving up to 40% cost reductions and 20-50% decreases in operational expenses through AI-driven automation (Databricks, 2025c). Additionally, optimized data pipelines have shown latency reductions exceeding 90% in banking implementations (Databricks, 2025a). For mid-sized companies (500-5,000 employees), structured intake prevents failures through registries and assessment frameworks, integrating with previous pillars (standards, catalogue, models, stewards) to create end-to-end AI governance (Databricks, 2025d; MLflow, 2025).

Why AI Use Cases Orchestration Matters

AI projects launch without data readiness assessments, leading to failures and compliance issues in mid-sized

enterprises. Research indicates 74% of companies struggle to scale AI value (Boston Consulting Group, 2024), with poor data quality being a critical factor. Organizations suffer income losses ranging from \$10 to \$14 million annually due to poor data quality (Deloitte Team, 2025), with additional studies showing average annual losses of \$12.9 million (Gartner, 2024b). Without structured orchestration, AI systems can experience hallucination rates of 15-30% for smaller or inadequately governed models (AllAboutAI Research Team, 2025), alongside model drift and compliance violations.

Pattern-based orchestration provides several critical benefits:

- **Risk Mitigation:** Data readiness assessments identify gaps before development begins, reducing project failures by 40-60% through proactive remediation (OvalEdge, 2025).
- **Risk Mitigation:** Data readiness assessments identify gaps before development begins, significantly reducing project failures. Organizations conducting thorough readiness assessments before AI initiatives report 62% higher success rates and achieve positive ROI up to 40% faster (Datasumi, 2025), while proper data readiness can prevent up to 60% of AI project abandonment (Gartner, 2025).
- **Efficiency Gains:** Model optimization techniques significantly improve performance. Case studies demonstrate latency reductions of 40-60% through quantization approaches (Markaicode, 2025), while optimized data processing frameworks achieve up to 59% reduction in response latency (Deepchecks, 2025). Additionally, efficient data preparation can reduce dataset size by 60% while accelerating training time by 3x (NVIDIA, 2025).

Correspondence concerning this article should be addressed to Gregory David Spehar, GiDanc AI LLC. Contact: greg@gidanc.ai

Acknowledgments: Thanks to the sources and best practices in data governance for mid-sized organizations.

Declaration of Interests: The author declares no competing financial interests.

- **Regulatory Compliance:** Integrated BCBS 239, GDPR, and ethical frameworks ensure AI deployments meet regulatory requirements from inception (Secoda Team, 2025).

Case studies demonstrate significant benefits from structured AI implementation in financial services. Organizations have achieved substantial operational improvements, including a 50% reduction in IT service calls and 19% increase in bank earnings (Nufar Gaspar, 2025), while financial institutions report average productivity gains of 20% across various functions (Bain & Company, 2024). AI-driven initiatives have delivered ROI within 18 months, with one case study showing \$35 million in fraud loss savings over this timeframe (Nufar Gaspar, 2025).

Business Problem and Process Flow

The fundamental challenge is that AI initiatives launch reactively without consideration for data preparedness, governance requirements, or ethical implications. Business units pursue AI projects independently, discovering data quality issues only after significant investment. This creates technical debt, compliance exposure, and organizational frustration when promising AI concepts fail due to preventable data problems.

The orchestrated process flow addresses this through systematic evaluation at each stage as shown in Figure 1.

This flow supports federated environments by centralizing strategic approvals (Council) while allowing domain-specific provisioning and development, enabling organizations to balance innovation velocity with governance rigor — an approach suited for rapidly scaling enterprises (Atlan Team, 2025).

Detailed Process Steps

The following Table 1 outlines proposed steps for AI use case orchestration, synthesizing best practices from governance frameworks (LightsOnData, 2025) and enterprise AI implementation patterns. The process integrates regulatory elements like GDPR in assessments and ensures council oversight for strategic alignment.

These steps enable structured AI deployment through pattern-based data preparation supporting various processing modes including batch ML workflows, streaming data pipelines, and real-time serving endpoints (Databricks, 2025b; Microsoft Learn, 2025).

Key Artifacts Produced

The following artifacts are produced during AI use case orchestration.

AI Use Case Intake Form Template

1 AI USE CASE INTAKE FORM

2

3 BUSINESS INFORMATION:

4 - Owner: [Name, Email]

5 - Department: [Sales/Finance/Operations]

6 - Priority: [High/Medium/Low]

7

8 PROBLEM STATEMENT:

9 - What business problem are we solving?

10 - What are current limitations?

11 - What is the expected business impact?

12

13 AI APPROACH:

14 - Type: [LLM/RAG, Predictive ML, Computer Vision]

15 - What data is needed?

16 - What is the timeline?

17

18 SUCCESS CRITERIA:

19 - How will we measure success?

20 - What performance level is required?

21

22 SUBMITTED BY: [Name] DATE: [Date]

Data Readiness Assessment Template

1 DATA READINESS ASSESSMENT

2

3 Use Case: [Name]

4 Assessed By: [DG Manager + Data Architect]

5 Date: [Date]

6

7 OVERALL READINESS: GREEN / YELLOW / RED

8

9 REQUIRED DATA SOURCES:

10 Source 1: [customers.gold_profile]

11 - Quality Score: 92%

12 - Gaps: Missing 'customer_segment' field

13 - Remediation Effort: 2 weeks

14

15 Source 2: [transactions.silver_history]

16 - Quality Score: 87%

17 - Gaps: 15% records missing timestamps

18 - Remediation Effort: 1 week

19

20 SUMMARY:

21 - Available Now: [2 sources]

22 - Needs Improvement: [1 source]

23 - Not Available: [0 sources]

24

25 ESTIMATED EFFORT: 3 weeks

26 DECISION: APPROVE with conditions / DEFER / REJECT

AI Use Case Registry Schema

```
1 CREATE TABLE governance.use_cases.ai_registry (  
2     use_case_id STRING,  
3     name STRING,  
4     owner STRING,  
5     ai_type STRING, -- 'LLM', 'ML_Predictive', 'Computer_Vision'  
6     status STRING, -- 'proposed', 'approved', 'in_dev', 'production'  
7     objective STRING,  
8     success_metrics STRING,  
9     data_sources ARRAY<STRING>,  
10    freshness_requirement STRING, -- 'real-time', 'daily', 'weekly'  
11    quality_threshold DECIMAL(3,2),  
12    sensitivity STRING, -- 'public', 'internal', 'confidential'  
13    pii_usage BOOLEAN,  
14    regulatory_requirements STRING, -- 'GDPR', 'BCBS_239', 'CCPA'  
15    approved_by STRING,  
16    approved_date DATE,  
17    assigned_engineer STRING,  
18    model_uri STRING,  
19    created_date DATE,  
20    last_updated TIMESTAMP  
21 );
```

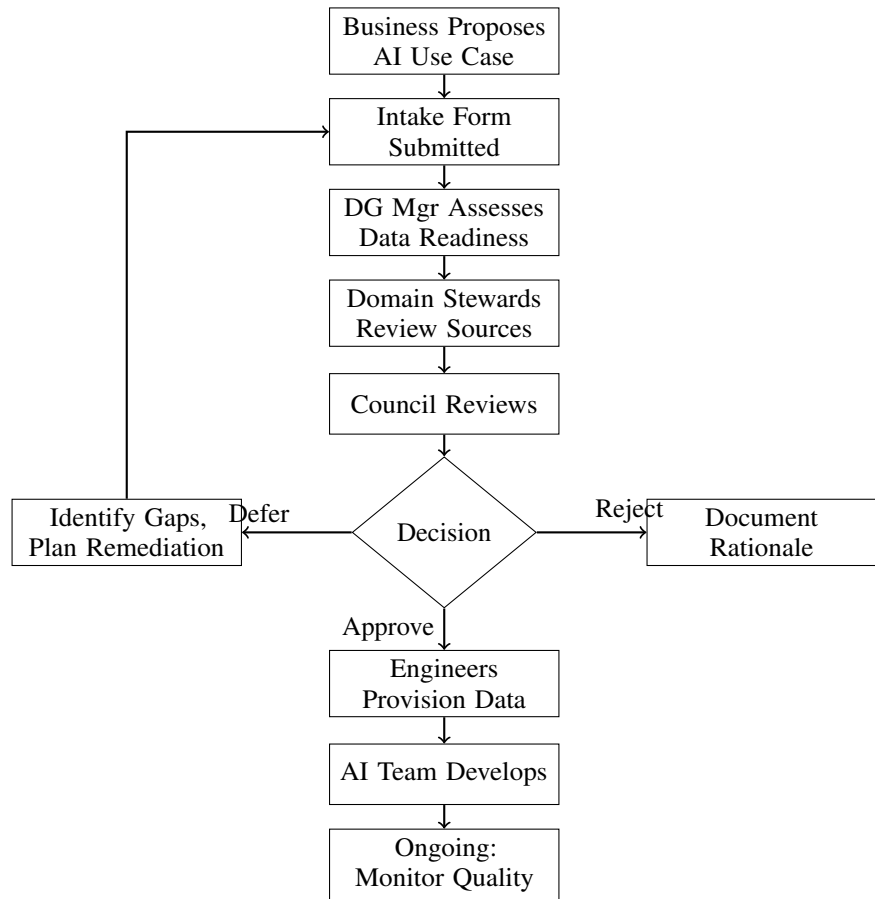


Figure 1

AI Use Case Orchestration Process Flow

AI Quality Monitoring Dashboard

```

1 CREATE LIVE TABLE ai_data_quality AS
2 SELECT
3     u.use_case_id,
4     u.name,
5     q.pass_rate as current_quality,
6     CASE
7         WHEN q.pass_rate >= u.quality_threshold
8         THEN 'HEALTHY'
9         ELSE 'CRITICAL'
10    END as status
11 FROM governance.use_cases.ai_registry u
12 JOIN LIVE.quality_dashboard q
13 ON q.dataset_name IN UNNEST(u.data_sources)
14 WHERE u.status = 'production';
    
```

These artifacts enable tailored orchestration patterns based on AI use case requirements. Organizations implement different data preparation strategies — from daily refresh cycles for LLM applications requiring current context, to batch processing for traditional ML workflows, to streaming architectures for real-time inference needs — guided by governance frameworks that balance flexibility with control (Alation Team, 2025).

Hypothetical Application: Implementation Example

To illustrate AI use case orchestration in practice, consider InsureTech Pro, a mid-sized insurance company with 2,800 employees offering property, auto, and life insurance products. This composite case draws from documented implementations in financial services and insurance sectors (Boston Consulting Group, 2024; Bradley, 2025).

The Hypothetical Business Problem

InsureTech’s Chief Innovation Officer wanted to deploy AI-powered claims fraud detection to reduce the \$8.7M annual loss from fraudulent claims. The business team proposed using an LLM to analyze claim descriptions and detect suspicious patterns. They approached IT requesting immediate access to claims data, expecting a 2-week timeline to production.

Investigation revealed critical data gaps: claims data existed across 4 legacy systems with inconsistent formats, 34% of historical claims lacked adjuster notes (key fraud indicators), and sensitive PII wasn’t properly governed. The business team had no understanding of AI data

Table 1*Detailed Process Steps for AI Use Case Orchestration*

Step	Role	Action	When	Output/Artifact
Use Case Proposal	Business Owner	Submit form with problem, AI type, metrics	Ongoing	AI Use Case Intake Form
Data Readiness Assessment	DG Manager + Architect	Review sources, quality, gaps	Within 1 week	Readiness Report (Red/Yellow/Green)
Domain Steward Review	Domain Stewards	Check compliance, privacy, feasibility	Next meeting	Approval with conditions
Council Review & Decision	Governance Council	Align with strategy, approve/defer/reject	Next quarterly	Decision Minutes
Use Case Registration	DG Manager	Add metadata for tracking	Within 1 week	Registry Entry
Data Provisioning	Data Engineer	Create schemas, transformations, access	2-4 weeks	AI-Ready Assets
Model Development	AI/ML Team	Train/register model	Weeks to months	Registered Model
Production Deployment	AI Team + Engineers	Deploy with monitoring	After testing	Deployment Docs
Ongoing Monitoring	Manager + Steward	Track quality, performance, compliance	Continuous	AI Quality Dashboard

Table 2*Key Artifacts for AI Use Case Orchestration*

Artifact	Description
AI Use Case Intake Form	Structured template capturing business problem, AI approach, data requirements, success criteria
Data Readiness Assessment	Evaluation framework scoring data sources (Red/Yellow/Green) with gap remediation estimates
AI Use Case Registry Schema	SQL table tracking all AI initiatives with metadata on type, status, data sources, regulatory requirements
AI Quality Monitoring Dashboard	Real-time tracking of model performance, data quality, and compliance metrics

requirements—they assumed "we have claims data" meant it was AI-ready.

Without orchestrated intake, this initiative would have followed the typical pattern: 3 months of development discovering data issues, \$450K invested before recognizing fundamental problems, project cancellation with business blaming "IT couldn't deliver," and organizational reluctance to attempt future AI projects.

- **Data Needed:** Historical claims (5 years), adjuster notes, settlement outcomes

- **Timeline:** 3 months to production

- **Success Metrics:** 80% fraud detection accuracy, <2% false positive rate

Step 2: Data Readiness Assessment (Week 2). The DG Manager and Data Architect evaluated required sources:

Applying the Use Case Orchestration Process

Following the process outlined in Figure 1, InsureTech implemented structured orchestration:

Step 1: Use Case Proposal (Week 1). The Claims VP submitted an AI Use Case Intake Form specifying:

- **Problem:** Reduce \$8.7M annual fraud losses by 40%
- **AI Type:** LLM for text analysis of claim descriptions

```

1 DATA READINESS ASSESSMENT
2
3 OVERALL READINESS: YELLOW (Conditional Approval)
4
5 REQUIRED SOURCES:
6 1. claims.silver_master
7   Quality: 78% (YELLOW)
8   Gaps: 34% missing adjuster_notes
9   Remediation: 4 weeks to backfill notes
10
11 2. claims.legacy_settlement
12   Quality: 62% (RED)
13   Gaps: Inconsistent fraud_flag encoding
14   Remediation: 6 weeks to standardize
15
16 3. customer.gold_profile
17   Quality: 94% (GREEN)
18   Gaps: None
19   Remediation: 0 weeks
20
21 SUMMARY:
22 - 1 source ready (GREEN)
23 - 1 needs improvement (YELLOW)
24 - 1 critical gaps (RED)
25
26 ESTIMATED EFFORT: 8 weeks remediation
27 RECOMMENDATION: DEFER until data improved

```

The assessment revealed that while claims data existed, quality fell below the 90% threshold required for reliable LLM training. The team estimated 8 weeks of remediation before AI development could begin.

Step 3: Domain Steward Review (Week 3). The Claims Domain Steward reviewed compliance and feasibility:

- **Privacy Concern:** Claim descriptions contain PII (names, addresses) requiring GDPR-compliant anonymization before LLM processing
- **Regulatory:** Fraud detection models subject to insurance regulatory fairness requirements—decisions must be explainable
- **Feasibility:** Backfilling 34% missing adjuster notes requires manual review by claims team (120 staff-hours)

The steward recommended: (1) Implement PII redaction in data preparation, (2) Use explainable AI techniques rather than black-box LLM, (3) Allocate claims team resources for note backfill.

Step 4: Council Review & Decision (Week 4). The Governance Council (CIO, CFO, Chief Claims Officer, Compliance Director) reviewed:

- **Strategic Alignment:** Fraud reduction aligned with company priorities
- **Risk Assessment:** Regulatory exposure from unexplainable AI decisions
- **Resource Reality:** Claims team bandwidth limited due to peak season

Decision: DEFER for 3 months with conditions:

1. Complete data remediation (8 weeks)
2. Implement PII redaction pipeline (2 weeks)

3. Schedule note backfill during slower Q1 period (4 weeks)
4. Revise AI approach to explainable model rather than pure LLM

Rather than reject the use case or approve with hidden risks, the Council provided a clear remediation path.

Step 5: Use Case Registration (Week 5). The DG Manager registered the use case:

```

1 INSERT INTO governance.use_cases.ai_registry VALUES (
2   'UC-2025-017',
3   'Claims Fraud Detection',
4   'claims_vp@insuretech.com',
5   'ML_Predictive', -- Changed from LLM per Council
6   'deferred',
7   'Reduce fraud losses 40% ($3.5M savings)',
8   'accuracy>=0.80, false_positive<=0.02',
9   ARRAY('claims.silver_master', 'claims.legacy_settlement'),
10  'daily', -- Refresh requirement
11  0.90, -- Quality threshold
12  'confidential',
13  TRUE, -- Contains PII
14  'GDPR,Insurance_Reg_Fairness',
15  'governance_council',
16  '2025-10-15',
17  NULL, -- Engineer assigned after approval
18  NULL, -- Model URI after deployment
19  '2025-10-01',
20  CURRENT_TIMESTAMP()
21 );

```

Steps 6-9: Remediation & Implementation (Months 2-5). During the 3-month deferral:

- Data engineers standardized fraud flags and backfilled adjuster notes
- Privacy team implemented PII redaction using Named Entity Recognition
- Data quality improved from 78% to 93%
- Council re-reviewed and APPROVED the use case
- Engineers provisioned feature tables with daily refresh
- AI team developed explainable gradient boosting model (not pure LLM)
- Model achieved 84% fraud detection accuracy with 1.3% false positive rate
- Production deployment with continuous quality monitoring

Results and Lessons Learned

Hypothetically, nine months post-orchestration (including 3-month deferral), InsureTech achieved measurable success: fraud detection accuracy of 84% (exceeding 80% target), \$3.2M annual savings from prevented fraud, and zero regulatory compliance issues. The structured process prevented the typical \$450K waste on premature development.

However, the deferral initially frustrated business stakeholders who expected immediate AI deployment. Executive communication proved critical—the CIO explained:

"We're not saying no to AI; we're ensuring AI succeeds. Three months of preparation prevents three years of problems." This reframing transformed deferral from rejection to investment in success.

The assessment process also revealed organizational gaps. Business teams lacked understanding of AI data requirements, assuming "we have data" equaled "data is AI-ready." InsureTech added AI literacy training for business leaders, covering: data quality needs, regulatory considerations, and realistic timelines.

The federated model proved essential: central Council provided strategic oversight and regulatory guidance, while Claims stewards contributed domain expertise on fraud patterns and data nuances. Pure centralization would have missed business context; pure decentralization would have created compliance risk.

This case demonstrates that orchestration's primary value isn't bureaucratic control—it's preventing expensive failures through proactive assessment, providing clear paths forward when gaps exist, and ensuring AI initiatives succeed rather than just start quickly.

Implementation Options for Mid-Sized Companies

Mid-sized firms can implement AI use case orchestration using Unity Catalog registries within Databricks environments. Based on typical enterprise implementations, estimated costs range from \$50-100K including platform licensing, workflow automation, and training. Alternatives include Azure Machine Learning for Microsoft-centric organizations, or open-source MLflow for cost-sensitive deployments (estimated \$20-50K setup including cloud infrastructure and initial configuration) (Microsoft Learn, 2025; MLflow, 2025).

Implementation of AI governance frameworks in financial services requires balancing performance optimization with regulatory compliance. Organizations must address requirements such as BCBS 239 for risk data management (N-iX, 2024) while aligning with ethical AI frameworks including NIST AI Risk Management standards (Bradley, 2025).

Effecting Change and Recommendations

Change management is crucial for AI governance success. Research indicates that 80% of data and analytics governance initiatives fail due to lack of organizational urgency and strategic positioning (Gartner, 2024a). Recommended strategies include:

- **Awareness Building:** Involve stakeholders early in the governance process and provide comprehensive training and workshops to help teams understand the value and reduce resistance to change (Secoda Team, 2025).
- **Role Clarity:** Use RACI matrices to define reviewer responsibilities, clearly distinguishing business sponsors, technical assessors, and council decision-makers.

Provide training on data readiness evaluation techniques (Data Quality Pro Team, 2025).

- **Role Clarity:** Use RACI matrices to define reviewer responsibilities, clearly distinguishing business sponsors, technical assessors, and council decision-makers in the governance framework (IT Governance Docs, 2025).
- **Pilot Approach:** Implement orchestration for high-value use cases through focused pilot programs, measure success through tangible business outcomes, then scale iteratively based on demonstrated value (Analytics8, 2025).
- **Ongoing Support:** Embed orchestration in AI culture through success story sharing and highlighting governance achievements to build stakeholder trust (Semarchy, 2025). Implement automated governance tools including data catalogs and metadata management systems to reduce administrative burden (Atlan, 2025b), and monitor adoption via usage metrics tracking whether certified data sources and AI registries are gaining organizational traction (Select Star, 2025).

Success requires treating orchestration as enablement, not gatekeeping—positioning assessments as "ensuring your AI succeeds" rather than "blocking your AI idea."

Discussion

This AI use cases pillar represents the strategic orchestration layer that transforms the five-pillar framework from theoretical governance to practical AI enablement. However, successful implementation demands more than process definition—it requires cultural transformation where "AI readiness" becomes as important as "AI innovation," and organizations value prevented failures as much as launched projects.

When Orchestration Models Struggle

Federated orchestration excels in organizations with mature AI understanding and collaborative cultures, but faces challenges in specific contexts. Companies in highly competitive industries may resist approval processes perceived as "slowing innovation," demanding streamlined assessments that risk inadequate evaluation. Organizations below 1,500 employees may lack dedicated AI teams to justify formal orchestration, with lightweight checklists proving more practical despite lower rigor.

Executive impatience represents a critical barrier. Leaders seeing competitors deploy AI rapidly may pressure teams to "skip governance and move fast," undermining orchestration before benefits materialize. Without sustained executive commitment that "failed AI costs more than delayed AI," orchestration devolves into paperwork rather than risk management.

Hidden Complexities and Assessment Realities

The stated 1-week data readiness assessment assumes organizations have established data quality metrics and lineage understanding. Without these foundations, assessments can take from weeks to several months depending on complexity and organizational scope (Atlan, 2025a), requiring teams to manually inventory sources, evaluate quality, and map dependencies — timelines that frustrate business stakeholders expecting rapid turnaround.

Moreover, "data readiness" proves contextual. The same dataset might be adequate for exploratory analytics but insufficient for production AI. Assessors must understand AI-specific requirements: LLMs needing diverse training examples, ML models requiring balanced datasets, computer vision demanding labeled images. This demands assessor skill evolution from traditional data governance to AI-aware evaluation.

Deferral decisions also introduce political complexity. Business teams may perceive deferrals as IT obstructionism rather than risk mitigation. Organizations must carefully communicate that deferrals provide clear remediation paths, not permanent rejections—a cultural shift requiring sustained leadership messaging.

Comparative Analysis: When Lightweight Approaches Excel

Formal orchestration represents comprehensive governance, but lighter approaches may prove superior in specific contexts. Early-stage AI organizations still learning capabilities may benefit from experimental freedom before implementing rigorous processes. Highly technical companies with strong data culture may find engineering-led assessments adequate without formal council oversight.

The choice depends on: AI maturity (experimental vs. production-scale), regulatory environment (highly regulated vs. minimal oversight), failure tolerance (mission-critical vs. experimental), and organizational size (dedicated AI teams vs. generalist IT). Organizations should honestly assess whether orchestration overhead provides commensurate risk reduction.

Evolution and Future Considerations

The AI landscape is evolving toward autonomous agents and federated learning, introducing orchestration complexities beyond current frameworks. Organizations implementing orchestration today should architect for: multi-agent system governance (how do we approve AI systems that spawn sub-agents?), federated model evaluation (assessing models trained on distributed data), and real-time adaptation oversight (governing models that self-tune in production).

As regulations like the EU AI Act mandate use case registration and risk assessment, orchestration shifts from best practice to compliance requirement. Future implementations must support regulatory reporting, with registries providing audit trails demonstrating due diligence in AI deployment decisions.

Critical Success Factors Revisited

Beyond process implementation, three organizational factors prove critical: executive patience valuing prevented failures over launch velocity, business-IT partnership where assessments are collaborative discovery rather than technical gatekeeping, and realistic expectations that orchestration adds 2-4 weeks to AI timelines but prevents 6-12 month failures. Organizations lacking these foundations see orchestration circumvented through shadow AI or abandoned under pressure.

The transformation from ad-hoc AI experimentation to orchestrated deployment succeeds when organizations embrace "AI readiness" as strategic capability. This requires cultural change where teams celebrate deferred projects that avoided expensive failures, where business leaders understand data quality fundamentals, and where "slower to start, faster to succeed" becomes accepted wisdom rather than frustrating bureaucracy.

AI use case excellence emerges not from perfect initial assessments, but from iterative learning—each orchestrated project builds organizational understanding of AI data requirements, regulatory considerations, and realistic success factors.

Organizations that treat orchestration as a continuous improvement system achieve sustainable AI scalability. By implementing cyclical governance processes that perpetually reflect and iterate on lessons learned (Anthology, 2024), and treating AI as a living system with feedback loops and evaluation testing (Microsoft Data Science, 2025), organizations build long-term value rather than one-off project successes.

Future Research Directions

This position paper proposes AI use case orchestration as a systematic approach to AI governance for mid-sized enterprises. While the framework synthesizes established governance principles with emerging AI requirements, several research questions warrant empirical investigation to validate, refine, and extend this work.

Empirical Validation of Orchestration Benefits

The claims presented—including 40-60% reductions in AI project failures and 20-30% efficiency gains—require rigorous empirical validation through longitudinal studies. Future research should employ quasi-experimental designs comparing organizations implementing orchestration frameworks against control groups pursuing traditional ad-hoc AI initiatives. Key research questions include:

- **Failure Rate Impact:** Do orchestrated AI projects demonstrate statistically significant reductions in abandonment rates, cost overruns, or timeline delays compared to unorchestrated initiatives? What mechanisms drive these improvements?

- **Time-to-Value Metrics:** Does structured assessment paradoxically accelerate overall project delivery by preventing mid-stream failures? How do initial assessment delays (2-4 weeks) compare to avoided remediation efforts (6-12 months)?
- **ROI Measurement:** What are the true costs of orchestration implementation (governance overhead, tooling, staffing) versus benefits (prevented failures, compliance assurance, reusable infrastructure)? At what organizational scale does positive ROI emerge?
- **Data Quality Correlation:** Do organizations with higher data quality maturity experience shorter assessment cycles and higher approval rates? Can we quantify the relationship between data readiness scores and AI project success?

Methodologically, this research requires access to project portfolios across multiple organizations, ideally combining quantitative metrics (project outcomes, costs, timelines) with qualitative case studies exploring causal mechanisms. Industry partnerships with consulting firms or technology vendors could provide necessary sample sizes.

Organizational and Contextual Success Factors

This paper identifies organizational contexts where orchestration may struggle—competitive industries prioritizing speed, smaller firms lacking dedicated AI teams, cultures resistant to governance. Future research should systematically investigate these boundary conditions through comparative case studies:

- **Organizational Size Thresholds:** The proposed 1,500-employee threshold for formal orchestration requires empirical validation. What governance approaches prove optimal for organizations at different scales (e.g., 100-500, 500-1,500, 1,500-5,000, 5,000+ employees)?
- **Industry Variation:** Do highly regulated industries (financial services, healthcare) demonstrate higher orchestration success rates than less-regulated sectors? How do competitive dynamics influence governance adoption and effectiveness?
- **AI Maturity Stages:** Should organizations adopt different orchestration models as they progress from experimental AI initiatives to production-scale deployments? What governance evolution paths prove most effective?
- **Cultural Prerequisites:** What organizational culture characteristics predict orchestration success? Can we measure factors like data literacy, cross-functional collaboration, and executive patience that enable governance adoption?

Grounded theory approaches and multiple case study methodologies would illuminate these contextual factors, building mid-range theory about governance fit with organizational characteristics.

Process Optimization and Tool Development

The proposed orchestration process represents a synthesis of best practices but requires refinement through implementation research. Action research methodologies could explore process improvements:

- **Assessment Duration Optimization:** What assessment techniques minimize time-to-decision while maintaining quality? Can automated data quality scanning, AI-powered lineage discovery, or pre-built templates accelerate evaluations?
- **Council Effectiveness:** How should governance councils be structured (size, composition, meeting frequency) to balance thoroughness with responsiveness? What decision-making protocols optimize approval quality?
- **Automation Opportunities:** Which orchestration steps benefit most from automation (intake forms, data quality checks, compliance verification, registry updates)? What level of human oversight remains essential?
- **Scalability Limits:** As AI initiatives proliferate, can orchestration processes handle increasing volumes without becoming bottlenecks? What architectural approaches (tiered review, domain delegation, risk-based assessment depth) maintain quality at scale?

Design science research approaches could iteratively develop, test, and refine orchestration tools and processes, contributing both practical artifacts and theoretical insights about AI governance design principles.

Comparative Governance Model Studies

This paper advocates federated orchestration balancing central oversight with domain autonomy. Comparative research should evaluate alternative models:

- **Centralized vs. Federated vs. Decentralized:** Under what conditions does each model prove superior? Do highly technical organizations benefit from engineering-led decentralization while regulated industries require central control?
- **Lightweight vs. Comprehensive Processes:** Can simplified governance (checklists, self-service assessments) achieve adequate risk mitigation for certain contexts? What criteria determine appropriate rigor levels?

- **Industry-Specific Adaptations:** How should orchestration frameworks adapt to unique requirements in healthcare (HIPAA, clinical safety), financial services (SOX, BCBS 239), or manufacturing (operational technology, supply chain)?

Multi-site case studies comparing organizations adopting different governance models would build understanding of contingency factors influencing model selection and success.

Methodological Considerations

Future research faces methodological challenges inherent to studying emerging organizational practices:

- **Longitudinal Access:** Orchestration benefits may materialize over years, requiring sustained research relationships with implementing organizations.
- **Counterfactual Estimation:** Demonstrating prevented failures requires estimating what would have occurred without orchestration—challenging without randomized control trials (often infeasible in organizational settings).
- **Confounding Variables:** Organizations adopting orchestration may differ systematically from non-adopters (higher data maturity, stronger executive support), complicating causal attribution.
- **Measurement Validity:** Governance "success" proves multidimensional (prevented failures, compliance assurance, organizational learning, stakeholder satisfaction), requiring validated measurement instruments.

Mixed-methods approaches combining quantitative metrics with rich qualitative data offer the most promising path forward, triangulating evidence across multiple data sources.

Contributions to Theory and Practice

This research agenda promises contributions across multiple domains. For *information systems theory*, investigations would extend understanding of data governance effectiveness, examining how governance structures influence AI system success. For *organizational theory*, research would illuminate how organizations balance innovation velocity with risk management, contributing to ambidexterity literature. For *AI ethics and governance*, empirical work would ground abstract principles in organizational realities, demonstrating how responsible AI translates to practice.

Practically, validated orchestration frameworks would provide mid-sized enterprises with evidence-based implementation guidance, reducing costly trial-and-error learning. Industry consortia could develop standardized maturity models, assessment tools, and training materials, accelerating governance adoption across sectors.

The transformation from position paper to validated theory requires sustained empirical investigation. Organizations

implementing orchestration frameworks should participate in research partnerships, contributing implementation data while gaining insights from comparative analyses. Academic-industry collaboration represents the optimal path for advancing both theoretical understanding and practical effectiveness of AI governance.

References

- Alation Team. (2025). *Federated data governance explained* [Accessed October 2025]. <https://www.alation.com/blog/federated-data-governance-explained/>
- AllAboutAI Research Team. (2025). *AI hallucination report 2025: Which AI hallucinates the most?* [Accessed October 2025]. <https://www.allaboutai.com/resources/ai-statistics/ai-hallucinations/>
- Analytics8. (2025). *How to build an effective data governance strategy* [Accessed October 2025]. <https://www.analytics8.com/blog/how-to-build-an-effective-data-governance-program-and-strategy/>
- Anthology. (2024). *Data governance and continuous improvement: Lessons learned for best practice* [Accessed October 2025]. <https://www.anthology.com/blog/data-governance-and-continuous-improvement-lessons-learned-for-best-practice>
- Atlan. (2025a). *Data governance readiness assessment: Questionnaire & how-to guide* [Accessed October 2025]. <https://atlan.com/data-governance-readiness-assessment/>
- Atlan. (2025b). *What is a data governance framework? guide & examples* [Accessed October 2025]. <https://atlan.com/data-governance-framework/>
- Atlan Team. (2025). *Federated data governance: Ultimate guide for 2024* [Accessed October 2025]. <https://atlan.com/known/data-governance/federated-data-governance/>
- Bain & Company. (2024). *AI in financial services survey shows productivity gains across the board* [Accessed October 2025]. <https://www.bain.com/insights/ai-in-financial-services-survey-shows-productivity-gains-across-the-board/>
- Boston Consulting Group. (2024). *74% struggle to scale AI value* [Accessed October 2025]. <https://www.bcg.com/publications/2024/ai-adoption-in-2024>
- Bradley. (2025). *Global AI governance: Five key frameworks explained* [Accessed October 2025]. <https://www.bradley.com/insights/publications/2025/08/global-ai-governance-five-key-frameworks-explained>
- Data Quality Pro Team. (2025). *How to implement change management for data quality and data governance* [Accessed October 2025]. <https://www.dataqualitypro.com/blog/data-quality-governance-change-management>
- Databricks. (2025a). *Banco bradesco transforms their CRM and doubles conversion rate* [Accessed October 2025]. <https://www.databricks.com/customers/banco-bradesco/lakeflow-declarative-pipelines>

- Databricks. (2025b). *New delta live tables capabilities* [Accessed October 2025]. <https://www.databricks.com/blog/new-delta-live-tables-capabilities>
- Databricks. (2025c). *Shifting to financial intelligence: Financial services at the data + ai summit 2025* [Accessed October 2025]. <https://www.databricks.com/blog/shifting-financial-intelligence-financial-services-data-ai-summit-2025>
- Databricks. (2025d). *Unity catalog facilitates collaboration for ml use cases* [Accessed October 2025]. <https://www.databricks.com/blog/machine-learning-unity-catalog>
- Datsumi. (2025). *AI readiness assessment for successful organisational transformation* [Accessed October 2025]. <https://www.datsumi.com/ai-readiness-assessment-for-successful-organisational-transformation>
- Deepchecks. (2025). *How do response time and latency factor into llm evaluation?* [Accessed October 2025]. <https://www.deepchecks.com/question/response-time-latency-llm-evaluation/>
- Deloitte Team. (2025). *Quality management in data governance* [Accessed October 2025]. <https://www.deloitte.com/ce/en/services/consulting/perspectives/bg-quality-management-in-data-governance.html>
- Gartner. (2024a). *Gartner predicts 80% of d&a governance initiatives will fail by 2027* [Accessed October 2025]. <https://www.gartner.com/en/newsroom/press-releases/2024-02-28-gartner-predicts-80-percent-of-data-and-analytics-governance-initiatives-will-fail-by-2027-due-to-a-lack-of-a-real-or-manufactured-crisis>
- Gartner. (2024b). *How to improve your data quality* [Accessed October 2025]. <https://www.gartner.com/smarterwithgartner/how-to-improve-your-data-quality>
- Gartner. (2025). *Lack of ai-ready data puts ai projects at risk* [Accessed October 2025]. <https://www.gartner.com/en/newsroom/press-releases/2025-02-26-lack-of-ai-ready-data-puts-ai-projects-at-risk>
- IT Governance Docs. (2025). *Data governance roles & responsibilities matrix (raci) template* [Accessed October 2025]. <https://www.itgov-docs.com/blogs/data-governance/governance-roles-and-responsibilities-matrix-raci-template>
- LightsOnData. (2025). *5 topics to cover in data governance council meetings* [Accessed October 2025]. <https://www.lightsondata.com/5-topics-to-cover-in-data-governance-council-meetings/>
- Markaicode. (2025). *How to reduce llm inference latency below 100ms: 7 proven optimization techniques* [Accessed October 2025]. <https://markaicode.com/reduce-llm-inference-latency-100ms/>
- Microsoft Data Science. (2025). *The ai governance gambit: Scale your ai without making headlines* [Accessed October 2025]. <https://medium.com/data-science-at-microsoft/the-ai-governance-gambit-scale-your-ai-without-making-headlines-6a613a193264>
- Microsoft Learn. (2025). *Mlflow integration for end-to-end ml* [Accessed October 2025]. <https://learn.microsoft.com/en-us/azure/databricks/machine-learning/mlflow/>
- MLflow. (2025). *Mlflow documentation* [Accessed October 2025]. <https://mlflow.org/docs/latest/>
- N-iX. (2024). *Data governance in banking and finance* [Accessed October 2025]. <https://www.n-ix.com/data-governance-banking-finance/>
- Nufar Gaspar. (2025). *The state of ai adoption in fintech: From efficiency gains to growth engines* [Accessed October 2025]. <https://www.nufargaspar.com/post/the-state-of-ai-adoption-in-fintech-from-efficiency-gains-to-growth-engines>
- NVIDIA. (2025). *Mastering llm techniques: Text data processing* [Accessed October 2025]. <https://developer.nvidia.com/blog/mastering-llm-techniques-data-preprocessing/>
- OvalEdge. (2025). *Bcbs 239 compliance through data governance* [Accessed October 2025]. <https://www.ovaledge.com/blog/bcbs-239-compliance-via-data-governance>
- Secoda Team. (2025). *What is change management in data governance?* [Accessed October 2025]. <https://www.secoda.co/glossary/what-is-change-management-in-data-governance>
- Select Star. (2025). *Defining data governance metrics and kpis* [Accessed October 2025]. <https://www.selectstar.com/resources/data-governance-metrics-and-kpis>
- Semarchy. (2025). *How to measure success with data governance metrics* [Accessed October 2025]. <https://semarchy.com/blog/how-to-measure-success-with-data-governance-metrics/>