

The Impact of Artificial Intelligence on Enterprise Architecture: A BAR Framework for Strategic, Governed, and Human- Centered AI Adoption

Bhagyeshkumar Chokhawala¹[0009-0001-5146-832X]

¹Comcast, Philadelphia, Pennsylvania, USA
bhagyeshkumar_chokhawala@comcast.com

Abstract

Artificial intelligence is changing enterprise architecture from a discipline centered on relatively deterministic applications and infrastructure into one that must govern data-dependent, probabilistic, continuously learning, and increasingly autonomous capabilities. Yet many organizations respond to AI demand through disconnected pilots, tool-specific decisions, and governance added after implementation. This conceptual paper examines how AI changes the scope, artifacts, decision rights, and operating model of enterprise architecture. Building on a practitioner article and an integrative synthesis of enterprise architecture, dynamic capabilities, human-centered AI, AI engineering, and governance literature, the paper introduces the BAR framework: Business Value and Purpose, Architecture Readiness and Integration, and Responsible Adoption and Realization. BAR links outcome-led use-case selection to data and platform readiness, human-AI collaboration, lifecycle controls, and measurable benefits. The paper also proposes an enterprise AI reference architecture, an evidence-gated adoption lifecycle, and an illustrative maturity index for determining when an initiative should remain an experiment, proceed to a controlled pilot, or scale as an enterprise capability. An AI-assisted customer retention scenario demonstrates how enterprise architects can connect recommendations, human authority, authorized execution, transaction integrity, and audit evidence. The paper argues that the central impact of AI on enterprise architecture is not the addition of a new technology layer; it is the need to make architecture a continuous socio-technical decision system that balances innovation, interoperability, trust, and realized value.

Keywords: artificial intelligence, enterprise architecture, AI adoption, responsible AI, human-AI collaboration, AI governance, architecture readiness, value realization

1. Introduction

Artificial intelligence (AI) has moved from specialist analytics teams into nearly every enterprise function. Human resources uses AI for workforce support and talent processes; sales and marketing use it for personalization and content generation; operations use it for forecasting and decision support; and software teams use it for coding, testing, documentation, and incident analysis. This diffusion creates a new challenge for enterprise architecture (EA). The architect is no longer asked only to rationalize applications, define technology standards, or guide cloud migration. The architect must help determine where AI should be used, what enterprise capabilities must support it, how probabilistic behavior is controlled, and how business value is measured after deployment.

The practitioner article that motivates this paper argues that successful AI adoption requires more than enthusiasm for language models or isolated automation. It highlights strategy and readiness, data and infrastructure, talent, integration and roadmaps, governance and ethics, and human-AI collaboration as mutually dependent conditions (Chokhawala, 2025). These concerns are consistent with research showing that AI introduces distinctive managerial challenges because model behavior depends on data, context, feedback, and organizational use rather than code alone (Berente et al., 2021). AI systems also accumulate forms of technical debt that are difficult to detect through conventional application governance, including data dependencies, feedback loops, model entanglement, and silent performance degradation (Sculley et al., 2015).

The resulting problem is a gap between AI experimentation and enterprise realization. Local teams can rapidly acquire an AI service, build a chatbot, or create a prototype agent. Enterprise value, however, depends on authoritative data, process integration, identity and access management, operational resilience, human decision rights, measurable outcomes, and continuing assurance. When these elements are addressed separately, organizations create AI sprawl: duplicated tools, inconsistent controls, unclear ownership, weak reuse, and rising operational risk. Enterprise architecture is positioned to bridge these concerns because it connects strategy, business capabilities, information, applications, technology, governance, and transition roadmaps.

This paper addresses three research questions:

- RQ1. How does AI change the scope, artifacts, and decision responsibilities of enterprise architecture?
- RQ2. What architecture conditions are required to move AI initiatives from experimentation to scalable enterprise capabilities?
- RQ3. How can governance, human oversight, adoption, and benefit realization be integrated into the AI architecture lifecycle?

The paper makes four contributions. First, it explains AI impact across seven EA domains. Second, it introduces the BAR framework - Business Value and Purpose, Architecture Readiness and Integration, and Responsible Adoption and Realization. Third, it provides an enterprise AI reference architecture and evidence-gated lifecycle. Fourth, it derives propositions and metrics that can be evaluated through future case studies, surveys, and longitudinal research.

2. Conceptual Background

2.1 AI as an enterprise capability

AI is often treated as a product feature or model-selection problem. At enterprise scale, it is better understood as a capability assembled from business rules, data, models, knowledge, processes, human roles, integration services, infrastructure, and governance. This distinction matters because model accuracy alone does not determine organizational performance. A technically capable model may fail when data access is delayed, recommendations are not integrated into work, users do not trust the output, accountability is unclear, or operational monitoring is absent. Conversely, a modest model can create significant value when embedded into a well-designed process with clear decision rights and feedback.

Digital transformation research similarly emphasizes that technology creates value by changing organizational properties, processes, and value creation rather than through acquisition alone (Vial, 2019). Dynamic capabilities theory explains why organizations must sense opportunities and threats, seize selected opportunities, and transform resources and operating models over time (Teece, 2007). EA can operationalize these capabilities for AI by sensing

use cases and risks, seizing opportunities through target architectures and investment roadmaps, and transforming data, platforms, processes, skills, and governance.

2.2 The automation-augmentation tension

AI adoption is frequently framed as a choice between automation and human augmentation. Automation transfers a task to a machine, whereas augmentation combines machine capabilities with human judgment. Raisch and Krakowski (2021) argue that the two are interdependent rather than clean alternatives. Automation can create new forms of human work, and augmentation can later enable further automation. The architecture implication is that decision authority must be designed dynamically. High-volume, reversible, low-impact decisions may support greater automation; ambiguous or consequential decisions require stronger human control, explanation, and escalation.

Human-centered AI research supports designs that combine high levels of useful automation with meaningful human control, responsibility, and self-efficacy (Shneiderman, 2020). For EA, this expands architecture beyond systems and interfaces to include role design, cognitive workload, override mechanisms, consent, and the quality of human-AI interaction. These concerns are especially important for agentic AI because an agent may plan, call tools, retain state, and trigger external effects across multiple systems.

2.3 AI engineering and lifecycle risk

Traditional software is primarily specified through deterministic logic. AI behavior also depends on training data, prompts, retrieval sources, model versions, evaluation sets, usage context, and feedback. Industrial research identifies workflow changes and new engineering challenges across data collection, model development, integration, testing, deployment, monitoring, and team coordination (Amershi et al., 2019). The hidden technical debt of machine learning includes undeclared consumers, unstable dependencies, configuration debt, feedback loops, and entangled changes that can undermine reliability (Sculley et al., 2015).

This means architecture reviews cannot end at design approval. Model and prompt changes, knowledge-base updates, access-policy changes, user behavior, and distribution shifts can alter system performance after release. EA therefore needs continuous architecture evidence: lineage, evaluation results, policy decisions, telemetry, incident records, cost, model versions, and benefit metrics.

2.4 Governance standards and assurance

Several standards provide foundations for responsible enterprise AI. The NIST AI Risk Management Framework organizes AI risk activities through Govern, Map, Measure, and Manage and identifies trustworthiness characteristics such as validity, safety, resilience, transparency, explainability, privacy, and fairness (NIST, 2023). ISO/IEC 42001:2023 specifies requirements for an organizational AI management system, while ISO/IEC 23894:2023 provides AI risk management guidance, and ISO/IEC 5338:2023 addresses AI system lifecycle processes. ISO/IEC 38507:2022 focuses on governance implications for governing bodies. These standards complement enterprise architecture methods such as the TOGAF Standard, which provides a structured approach for developing, maintaining, and using enterprise architecture (The Open Group, 2022).

The standards define what organizations should govern, but architecture must translate them into operating controls and reusable design patterns. The BAR framework performs this translation by connecting business purpose, architecture readiness, and responsible realization at portfolio, solution, and operational levels.

3. Research Method

This study uses a conceptual framework-development method rather than an empirical design. The starting problem and practice constructs were extracted from the author's article on the impact of AI on enterprise architecture (Chokhawala, 2025). These constructs were compared and synthesized with peer-reviewed research on AI management, digital transformation, dynamic capabilities, automation and augmentation, human-centered AI, AI engineering, technical debt, ethics, and algorithmic auditing. They were also aligned with the NIST AI RMF, ISO/IEC AI governance and lifecycle standards, and the TOGAF Standard.

Framework construction followed three stages. First, recurring adoption concerns were grouped into business, architecture, and responsibility domains. Second, the domains were decomposed into decision areas, artifacts,

measures, and lifecycle gates. Third, internal coherence was examined through an illustrative customer-retention scenario in which AI supports recommendation and authorized action execution. Because the framework has not yet been empirically validated, the maturity weights, thresholds, and propositions presented below should be treated as testable design hypotheses rather than universal constants.

4. How AI Changes Enterprise Architecture

AI expands EA work in at least seven domains. Table 1 summarizes the shift from conventional concerns to AI-specific architecture responsibilities.

EA domain	Conventional emphasis	AI-era architecture responsibility
Strategy and portfolio	Align investments, capabilities, and roadmaps	Prioritize outcome-led AI use cases; manage model and vendor portfolios; distinguish experiment, product, and enterprise capability; define benefit owners and risk tiers
Business and process architecture	Map capabilities, value streams, roles, and processes	Redesign human-AI work; locate decision points; define autonomy, escalation, override, and accountability; prevent automation of defective processes
Data and information architecture	Define information ownership, models, quality, and exchange	Add training and inference lineage, retrieval sources, semantic context, vector data, data rights, quality thresholds, and feedback controls
Application and integration architecture	Define services, interfaces, and system responsibilities	Establish model gateways, agent tools, policy-mediated APIs, event orchestration, transaction integrity, and fallback behavior
Technology architecture	Standardize platforms, networks, compute, and runtime	Plan accelerators, cloud/edge placement, latency, model operations, availability, capacity, portability, and cost controls
Security, risk, and governance	Apply standards, security patterns, and architecture assurance	Classify AI risk, enforce identity and least privilege, protect prompts and data, validate models, record decisions, and manage incidents and drift
Operating model and talent	Define architecture roles, boards, standards, and communities	Create cross-functional AI governance, product-model-platform ownership, AI literacy, evaluation skills, change management, and continuous evidence review

Table 1. Impact of AI across enterprise architecture domains.

The most significant change is temporal. Conventional architecture commonly emphasizes target-state planning and design conformance before deployment. AI requires architecture to remain active during operation because behavior, cost, data, and risk can change. The architecture repository must therefore evolve from a catalog of applications and diagrams into a system of record for AI capabilities, datasets, models, prompts, knowledge sources, policies, owners, evaluations, dependencies, and evidence.

5. The BAR Framework

The BAR framework translates the article's call to "raise the bar" into three interdependent architecture conditions. Business Value and Purpose ensures that AI is applied to a meaningful outcome rather than adopted as a generic technology. Architecture Readiness and Integration ensures that the enterprise can support the use case reliably and reuse the resulting capabilities. Responsible Adoption and Realization ensures that people, controls, and measurement sustain value after deployment. None of the dimensions can compensate fully for failure in another. A strong model without trusted data and workflow integration will not scale; a robust platform without a valuable use case creates cost without benefit; governance without adoption and operational evidence becomes documentation rather than control.

BAR Framework for Enterprise AI Adoption

Raise the BAR: align value, readiness, and responsibility before scaling AI




Figure 1. BAR framework for enterprise AI adoption.

5.1 B - Business Value and Purpose

The business dimension begins with a specific problem, decision, user, and intended outcome. It asks why AI is appropriate compared with rules, analytics, process redesign, or conventional automation. Use cases should be associated with a capability, value stream, process owner, affected stakeholders, and measurable baseline. Architecture work includes identifying where AI changes a decision, what information it consumes, what action may follow, and who remains accountable.

An outcome-led use-case definition should specify at least: the business objective; current performance; target improvement; decision frequency and reversibility; affected users and customers; required data; expected human role; downstream actions; risk tier; and benefit owner. This makes it possible to compare AI investments using value, feasibility, risk, and strategic reuse rather than novelty.

5.2 A - Architecture Readiness and Integration

Architecture readiness concerns the enterprise conditions that determine whether a promising use case can operate at scale. Data must be authoritative, accessible, timely, lawful, and sufficiently representative. Platforms must support model access, evaluation, versioning, observability, security, resilience, and cost control. Integration must connect AI recommendations to workflows and systems through governed APIs, events, and orchestration rather than fragile point-to-point scripts. Knowledge sources and retrieval mechanisms require provenance and freshness controls.

Readiness also includes nonfunctional requirements. Response latency, availability, throughput, privacy, explainability, portability, recovery, and energy or compute constraints vary by use case. An architect should make these requirements explicit before tool selection. Build-versus-buy decisions should consider not only model capability but data residency, contractual terms, model-change notification, evaluation transparency, exit cost, and the ability to route workloads across models.

5.3 R - Responsible Adoption and Realization

Responsible adoption combines governance, human-centered design, talent, change, and ongoing benefit realization. Risk controls should be proportional to impact. Low-risk drafting support may require user guidance and data restrictions; customer-affecting recommendations require stronger evaluation and explanation; autonomous financial, service, or safety actions require explicit authorization, transaction limits, reversibility, and escalation. Internal algorithmic auditing research emphasizes embedding accountability activities throughout the lifecycle rather than evaluating harm only after deployment (Raji et al., 2020).

Adoption is equally important. Users need clear expectations about what the system can and cannot do, how to challenge an output, when to escalate, and how their feedback is used. Training must cover implementers, users,

product owners, risk teams, and executives. Benefit realization requires measurement after launch: adoption, decision quality, cycle time, customer outcomes, cost, incidents, overrides, and model performance. Ethics principles such as transparency, fairness, privacy, responsibility, and non-maleficence become meaningful only when translated into roles, controls, evidence, and corrective action (Jobin et al., 2019).

5.4 Illustrative BAR maturity index

For portfolio triage, the three dimensions can be scored on a five-point scale. The following illustrative index gives slightly greater weight to value and readiness while preserving responsibility as a near-equal condition:

$$\text{BAR Score} = 0.35(\text{Business Value}) + 0.35(\text{Architecture Readiness}) + 0.30(\text{Responsible Adoption})$$

A score below 2.0 indicates that the initiative should remain in discovery or sandbox experimentation. Scores from 2.0 to 2.9 support a constrained pilot with explicit learning objectives. Scores from 3.0 to 3.9 may support controlled production scale when critical risk dimensions meet minimum thresholds. Scores of 4.0 or higher indicate stronger enterprise-scale readiness. A minimum score should also be required for each dimension so that high expected value cannot mask unacceptable risk or weak architecture.

BAR dimension	Assessment considerations	Minimum evidence
Business Value and Purpose	Problem clarity; strategic alignment; measurable baseline; accountable owner; user and customer value; reuse potential	Outcome metrics, benefit hypothesis, capability map, use-case charter
Architecture Readiness and Integration	Data quality; integration maturity; platform capability; security; resilience; observability; portability; lifecycle cost	Data product assessment, target architecture, NFRs, dependency map, cost model
Responsible Adoption and Realization	Risk classification; human authority; explainability; privacy; training; adoption; audit evidence; monitoring; corrective action	Impact assessment, control plan, RACI, evaluation report, operating playbook

Table 2. BAR assessment dimensions and evidence.

6. Enterprise AI Reference Architecture

Figure 2 provides a vendor-neutral reference architecture. The model separates experience and work, business capabilities and processes, AI services and agents, integration, data and knowledge, and platform infrastructure. Governance and operational evidence are cross-cutting because they apply to every layer. This prevents AI governance from becoming a separate review document disconnected from design and runtime behavior.




Figure 2. Vendor-neutral enterprise AI reference architecture.

6.1 Experience and work layer

The top layer includes customer channels, employee tools, partner experiences, and human decision points. It defines disclosure, consent, accessibility, interaction design, and fallback experiences. The architecture must specify whether AI drafts, recommends, decides, or acts and how the user understands that role.

6.2 Business capability and process layer

This layer anchors AI to business capabilities, value streams, policies, and outcomes. It prevents tool-driven architecture by locating AI within an owned process and clarifying the consequences of a recommendation or action. Business rules, decision models, BPM, and human approvals coexist with AI rather than being replaced automatically.

6.3 Agent and AI service layer

Reusable AI services include model gateways, prompt and policy templates, retrieval, evaluation, content safety, agent orchestration, memory, and tool catalogs. A gateway can centralize approved models, identity, routing, quotas, logging, and policy. Agent tools should expose outcome-oriented APIs with narrow authorization instead of direct database or unrestricted system access.

6.4 Integration and orchestration layer

APIs, events, workflow, and process orchestration connect AI to enterprise transactions. This layer is essential when AI recommendations lead to service changes, payments, orders, provisioning, or other state changes. Transaction integrity requires idempotency, validation, authorization, state tracking, compensating actions, and clear separation between recommendation and execution.

6.5 Data, knowledge, platform, and evidence layers

The data and knowledge layer includes authoritative data products, metadata, lineage, vector stores, ontologies, and knowledge graphs. The platform layer provides compute, storage, runtime, model operations, network, edge capabilities, and resilience. Cross-cutting governance controls identity, privacy, security, risk, policy, and human

authority. Cross-cutting evidence includes evaluation, quality, drift, cost, audit trails, incidents, and realized outcomes. Together, these layers make trust observable rather than assumed.

6.6 Architecture principles

- Purpose before model: select AI only after defining the outcome, user, decision, and alternative solutions.
- Capability over tool: build reusable enterprise capabilities rather than independent vendor-specific features.
- Governed data access: enforce ownership, lineage, minimization, quality, and purpose-based access.
- Human authority proportional to impact: increase review, explanation, approval, and escalation as consequences rise.
- Reversibility and safe failure: provide fallback, bounded actions, idempotency, and compensation for external effects.
- Evidence by design: generate evaluation, policy, version, decision, and outcome evidence as part of operation.
- Portability and controlled dependency: manage model, data, platform, and vendor lock-in explicitly.
- Continuous validation: monitor performance, drift, adoption, incidents, and benefits throughout the lifecycle.

7. Evidence-Gated Adoption Lifecycle

A roadmap should increase investment and autonomy only when evidence supports the next step. Figure 3 presents six stages. Discovery establishes the problem and alternatives. Qualification tests value, ownership, risk, and readiness. Design defines the target architecture and controls. Pilot produces technical, operational, user, and risk evidence. Scale establishes reusable platform and operating capabilities. Operate and Learn uses telemetry and outcomes to improve or retire the system.




Figure 3. Evidence-gated enterprise AI adoption lifecycle.

Lifecycle stage	Architecture gate	Required evidence
Discover	Is there a material problem and is AI an appropriate option?	Problem statement, baseline, affected users, alternatives, preliminary value
Qualify	Should the use case enter the AI portfolio?	BAR assessment, accountable owner, risk tier, data feasibility, benefit hypothesis
Design	Is the target solution coherent and governable?	Reference architecture mapping, NFRs, human authority, controls, evaluation plan, cost model
Pilot	Does evidence support controlled production use?	Offline and live evaluation, security testing, usability, adoption, error analysis, operating readiness
Scale	Can the capability operate reliably and be reused?	SLOs, support model, training, capacity, observability, disaster recovery, standardized components

Lifecycle stage	Architecture gate	Required evidence
Operate and Learn	Is the system producing sustained value within risk tolerance?	Telemetry, drift, incidents, override rates, outcome metrics, cost, corrective action, retirement decision

Table 3. Architecture gates for AI adoption.

8. Illustrative Enterprise Scenario: AI-Assisted Customer Retention

Consider a customer who contacts a service representative to discuss concerns and possible alternatives before the service is disconnected. In a conventional process, the representative must listen, search multiple systems, interpret eligibility rules, compare offers, explain options, and initiate execution. The practitioner article proposes using AI to quickly gather options while preserving the representative's role in empathy, critical thinking, and relationship-building (Chokhawala, 2025).

In the BAR design, the business objective is not to deploy a conversational model. It is to improve retention, decision quality, customer understanding, and representative efficiency without increasing harmful or unauthorized actions. The architecture separates five responsibilities:

- Context assembly: retrieve authorized customer, service, interaction, policy, and eligibility information with provenance.
- AI reasoning support: summarize concerns, identify relevant options, explain trade-offs, and estimate likely outcomes.
- Human decision and communication: the representative validates the context, applies empathy and judgment, discusses options, and receives customer agreement.
- Authorized execution: an automation service invokes outcome-oriented APIs only after identity, consent, policy, and transaction validation.
- Evidence and learning: the system records sources, model version, recommendation, human changes, consent, execution result, customer outcome, and feedback.

This design avoids two extremes. A fully manual process underuses AI and imposes search and documentation burdens on the representative. A fully autonomous interaction can reduce empathy, provide generic responses, or execute an unsuitable action. The architecture instead combines AI speed and scale with human responsibility and controlled automation.

Scenario risk	Architecture controls
Incorrect or incomplete recommendation	Grounding in authoritative sources; eligibility rules; confidence threshold; human validation; fallback search
Unauthorized service change	Strong identity, explicit consent, policy check, scoped token, action limits, separation of recommendation and execution
Inconsistent multi-system state	Outcome API, orchestration, idempotency, saga/compensation, reconciliation, transaction audit
Bias or unfair offer treatment	Segment evaluation, fairness review, policy constraints, explanation, monitoring by outcome group
Sensitive-data leakage	Data minimization, retrieval filters, encryption, redaction, approved model routing, retention controls
Low workforce adoption	Representative co-design, training, transparent reasoning, override without penalty, feedback loop, workload measurement

Table 4. Risks and controls for the illustrative customer-retention scenario.

Evaluation should combine AI, process, customer, and operational measures: recommendation relevance; unsupported-claim rate; representative acceptance and modification; average handling time; first-contact resolution; retention; customer satisfaction; execution failure; compensation rate; policy exceptions; incidents; and total cost per resolved case. The architecture is successful only when the combined process improves, not merely when the model produces plausible text.

9. Research Propositions and Evaluation Model

The conceptual synthesis supports five propositions that can guide empirical research:

- P1.** Business-value alignment positively influences realized AI performance because use cases are anchored to measurable outcomes, accountable owners, and process change.
- P2.** Architecture readiness positively moderates the relationship between AI capability and enterprise value; model capability creates less value when data, integration, platform, or operational readiness is weak.
- P3.** Responsible governance increases the probability of sustainable scaling by improving trust, accountability, risk detection, and organizational permission to reuse AI capabilities.
- P4.** Human-centered augmentation improves adoption and decision quality when human authority, explanation, workload, and escalation are designed explicitly.
- P5.** Continuous architecture evidence reduces AI technical debt and residual risk by enabling detection of drift, dependency changes, policy violations, incidents, and benefit erosion.

A complementary value-realization model can be expressed conceptually as follows:

$$\text{Realized AI Value} = \text{Expected Benefit} \times \text{Adoption} \times \text{Reliability} \times \text{Trust} - \text{Lifecycle Cost} - \text{Residual Risk Exposure}$$

The multiplicative terms reflect an important architecture insight: benefit collapses when adoption, reliability, or trust approaches zero. Lifecycle cost includes data, integration, evaluation, platform, model usage, monitoring, support, change, and retirement. Residual risk exposure represents expected loss after controls. Future studies can operationalize these constructs and test whether BAR maturity predicts scale, performance, incident rates, and benefit persistence.

10. Discussion

10.1 Enterprise architecture as a dynamic AI capability

The framework positions EA as a dynamic capability rather than a documentation function. Sensing occurs through AI opportunity discovery, technology scanning, risk intelligence, and capability diagnostics. Seizing occurs through portfolio decisions, target architectures, reusable platform choices, and staged investment. Transforming occurs through process redesign, data-product development, operating-model change, skills, governance, and continuous learning. This view is consistent with dynamic capabilities theory and with the need for architecture to adapt as AI technologies and organizational conditions evolve.

10.2 From architecture review to continuous assurance

AI weakens the assumption that an approved design remains stable after deployment. Enterprise architects therefore need mechanisms for continuous assurance. Architecture policies should be machine-enforceable where possible: model access through gateways, scoped agent tools, data classification, mandatory telemetry, evaluation thresholds, and version registration. Architecture boards can then focus on exceptions, trade-offs, and evidence rather than manually checking every routine control.

10.3 Changing roles and competencies

The EA role expands in five directions: AI capability strategist, socio-technical process designer, data and platform steward, policy and control architect, and value-realization analyst. Architects need enough AI literacy to understand model behavior, evaluation, retrieval, agents, and model operations, but they do not replace data scientists or engineers. Their distinctive responsibility is to integrate these specialties into a coherent enterprise system and to make cross-domain trade-offs visible.

10.4 Interoperability and vendor strategy

Rapid AI innovation increases the risk of tool fragmentation and lock-in. A model gateway, reusable evaluation service, common identity, governed tool interfaces, standard telemetry, and portable knowledge assets can isolate business capabilities from specific models. Interoperability should include semantic and policy consistency, not only technical connectivity. An agent that can call an API is not interoperable in an enterprise sense unless it understands the outcome contract, authorization, data meaning, failure behavior, and evidence obligations.

10.5 Ethical and organizational implications

AI architecture choices redistribute authority, visibility, workload, and opportunity. Governance must therefore involve business owners, users, risk, legal, security, data, technology, and affected stakeholders. High-level principles are necessary but insufficient; they must be translated into lifecycle decisions and operational evidence. The NIST and ISO frameworks provide strong foundations, while the BAR framework focuses on their integration with enterprise value, architecture, and adoption.

11. Limitations and Future Research

The paper is conceptual and uses an illustrative scenario rather than empirical data. The BAR dimensions and weights require validation across organizations, industries, and AI use-case types. Risk thresholds will differ in regulated, safety-critical, public-sector, and consumer contexts. The framework also abstracts from the detailed regulatory requirements of individual jurisdictions.

Future research should develop and validate a BAR assessment instrument; test the relationship between architecture maturity and realized AI value; compare centralized, federated, and platform-based AI operating models; study how architecture decision rights change with agentic autonomy; and examine whether continuous evidence reduces incidents and technical debt. Longitudinal case studies could measure how use cases move from pilot to scale and how governance and architecture practices evolve after failure, model change, or organizational restructuring.

12. Conclusion

AI has a broader impact on enterprise architecture than adding models, accelerators, or new cloud services. It changes the nature of enterprise capabilities, decision rights, integration, risk, operations, and human work. Organizations that treat AI as a collection of local tools may achieve rapid demonstrations but struggle with reuse, trust, cost, and production stability.

The BAR framework provides a practical conceptual response. Business Value and Purpose anchor AI to outcomes and ownership. Architecture Readiness and Integration create the data, platform, interoperability, security, and operational conditions for scale. Responsible Adoption and Realization embeds governance, human-AI collaboration, change, monitoring, and benefits. The reference architecture and evidence-gated lifecycle translate these dimensions into design and operating decisions. Enterprise architects can therefore help move AI from hype to strategy by making it a governed, measurable, and adaptive enterprise capability.

References

- Amershi, S., Begel, A., Bird, C., DeLine, R., Gall, H., Kamar, E., Nagappan, N., Nushi, B., & Zimmermann, T. (2019). Software engineering for machine learning: A case study. *Proceedings of the 41st International Conference on Software Engineering: Software Engineering in Practice*, 291-300. <https://doi.org/10.1109/ICSE-SEIP.2019.00042>
- Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial intelligence. *MIS Quarterly*, 45(3), 1433-1450. <https://doi.org/10.25300/MISQ/2021/16274>
- Chokhawala, B. (2025, May 2). From hype to strategy: Navigating AI integration in the modern enterprise. Medium. <https://medium.com/@bhagyesh.chokhawala/impact-of-ai-on-enterprise-architecture-7fde90656989>
- Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., Duan, Y., Dwivedi, R., Edwards, J., Eirug, A., Galanos, V., Ilavarasan, P. V., Janssen, M., Jones, P., Kar, A. K., Kizgin, H., Kronemann, B., Lal, B., Lucini, B., ... Williams, M. D. (2021). Artificial intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*, 57, 101994. <https://doi.org/10.1016/j.ijinfomgt.2019.08.002>
- International Organization for Standardization. (2022). ISO/IEC 38507:2022, Information technology - Governance of IT - Governance implications of the use of artificial intelligence by organizations.
- International Organization for Standardization. (2023a). ISO/IEC 23894:2023, Information technology - Artificial intelligence - Guidance on risk management.
- International Organization for Standardization. (2023b). ISO/IEC 42001:2023, Information technology - Artificial intelligence - Management system.
- International Organization for Standardization. (2023c). ISO/IEC 5338:2023, Information technology - Artificial intelligence - AI system life cycle processes.
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1, 389-399. <https://doi.org/10.1038/s42256-019-0088-2>
- National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). U.S. Department of Commerce. <https://doi.org/10.6028/NIST.AI.100-1>
- National Institute of Standards and Technology. (2024). Artificial intelligence risk management framework: Generative artificial intelligence profile (NIST AI 600-1). U.S. Department of Commerce. <https://doi.org/10.6028/NIST.AI.600-1>
- Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation-augmentation paradox. *Academy of Management Review*, 46(1), 192-210. <https://doi.org/10.5465/amr.2018.0072>
- Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 33-44. <https://doi.org/10.1145/3351095.3372873>
- Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe and trustworthy. *International Journal of Human-Computer Interaction*, 36(6), 495-504. <https://doi.org/10.1080/10447318.2020.1741118>
- Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. *Advances in Neural Information Processing Systems*, 28, 2503-2511.
- Teece, D. J. (2007). Explicating dynamic capabilities: The nature and microfoundations of sustainable enterprise performance. *Strategic Management Journal*, 28(13), 1319-1350. <https://doi.org/10.1002/smj.640>
- The Open Group. (2022). The TOGAF Standard (10th ed.). The Open Group.
- Vial, G. (2019). Understanding digital transformation: A review and a research agenda. *The Journal of Strategic Information Systems*, 28(2), 118-144. <https://doi.org/10.1016/j.jsis.2019.01.003>