

# A Governance Centric Conceptual Framework for Enterprise AI Transformation Across Intelligent Workflow Platforms

Girish Rajagopalan Sessa  
Cognizant Technology Solutions

## Abstract

Enterprise AI transformation requires coordinated strategy, readiness, governance, architecture, implementation, operations, and improvement. This paper proposes the vendor-neutral Enterprise AI Transformation Framework (EATF), integrating these seven phases with eight supporting reference models. Literature-informed conceptual synthesis, embedded architecture views, and an illustrative enterprise scenario explain its application. The proposed conceptual innovation is an evidence-linked enterprise decision architecture connecting business intent, readiness, risk controls, technical release evidence, and value realization; effectiveness and assessment metrics remain unvalidated. EATF complements enterprise architecture, AI risk management, and MLOps rather than replacing them.

*Keywords—Enterprise AI; governance; readiness; workflow platforms; MLOps; agentic AI.*

## I INTRODUCTION AND RELATED WORK

Digital transformation depends on organizational change as well as technology [1]. AI readiness and capability research highlights strategy, data, skills, resources, and culture [2], [3]. Ethics guidelines and algorithmic auditing address responsible use and accountability [4], [5], while MLOps connects model development to production operations [6]. These complementary perspectives motivate integration; they do not establish that enterprise-wide frameworks are absent. EATF links their requirements to a shared transformation lifecycle and workflow architecture.

### A Problem statement

Enterprise AI initiatives can satisfy model-development checks while remaining disconnected from business authorization, organizational readiness, cross-system accountability, and benefit realization. The specific problem addressed by EATF is the absence of an explicit coordination mechanism in the approaches compared here for carrying evidence and decision authority from an approved business use case through deployment, operation, and portfolio review. Without that connection, teams can duplicate assurance work, leave ownership or approval conditions unclear, and evaluate technical performance separately from workflow outcomes. These are the problems the proposed framework is designed to address, not measured findings of this study.


## B Research gap and intended contribution

The gap identified through the literature-informed comparison is one of coordination across complementary disciplines. Organizational AI-readiness and capability studies explain prerequisites for adoption [2], [3]; auditing research addresses accountability evidence [5]; MLOps connects engineering and operations [6]; and TOGAF, NIST AI RMF, and ISO/IEC 42001 provide architecture, risk-management, and management-system foundations [16]–[18]. These sources support essential parts of transformation, but the comparison in this paper does not specify a single shared design that connects readiness remediation, business authorization, technical release approval, residual-risk acceptance, and portfolio continuation through the same use-case evidence record.

EATF fills this identified design gap by proposing an evidence-linked enterprise decision architecture: seven transformation phases and eight supporting models coordinate a persistent use-case record, distinct but connected decision rights, and feedback from incidents and realized benefits. The research question is whether this coordination improves traceability and accountable decisions while reducing duplicated assurance effort, without excessive approval overhead. The contribution is a testable integration pattern; a systematic comparative review is still required to establish its novelty relative to other enterprise-wide approaches, and empirical studies are required to establish effectiveness.

## II RESEARCH METHODOLOGY

Figure 1 expresses the manuscript's qualitative synthesis process, informed by design-science principles [7], [8]. Concepts from the cited literature and frameworks are mapped to EATF responsibilities. In this revision, the mappings were examined through an internal conceptual consistency review and a lifecycle coverage review, as described below. These checks establish design plausibility within the stated scope; they do not constitute independent expert validation, a systematic literature review, an ISO conformity assessment, or empirical scalability testing.



**Fig. 1. Research methodology for conceptual development and proposed validation.**

### A Mapping decisions and validation status

Each mapping was reviewed against three criteria: whether the source framework addresses the stated responsibility, whether the assigned EATF phase or supporting model has a corresponding purpose, and whether the claimed relationship exceeds what the source establishes. The review used the framework descriptions and references [6], [16]–[18], with many-to-many relationships retained where responsibilities cross phases. For example, NIST Govern maps to EAGM throughout the lifecycle; Map informs strategy, readiness, and architecture; Measure informs EAWRF testing and operations; and Manage informs risk treatment, incident response, and improvement. Governance was not assigned exclusively to one phase, and MLOps was not represented as a single standardized framework.

The resulting decisions are documented in Table I and Section V as conceptual correspondences. Where EATF adds a shared register, coordinated decision rights, or portfolio feedback, these are identified as proposed design mechanisms rather than requirements attributed to NIST, ISO, or MLOps. No external reviewers, agreement statistics, interview findings, or validation sample are reported because none has been obtained. Future validation should use a mapping worksheet containing source provision, EATF responsibility, rationale, evidence artifact, and ambiguity; independent architecture, AI-risk, and engineering experts should review it, record disagreements, and reconcile revisions. Agreement and unresolved exceptions should be reported alongside the final mapping.

### B Completeness assessment

Completeness was assessed qualitatively against the declared seven-phase lifecycle and eight supporting models, rather than against every possible enterprise AI requirement. The coverage review checked that strategy has an investment owner and value hypothesis; readiness has capability gaps and remediation; governance has policies and decision authority; architecture has data and integration dependencies; implementation has testing and release evidence; operations has monitoring and incident response; and improvement has corrective work and value review. It also checked that the eight models have explicit roles and that evidence can pass between phases through the shared use-case record. These responsibilities are represented in Sections III–VI.

This assessment indicates coverage of the paper’s defined conceptual scope, not exhaustive completeness. The manuscript does not establish a clause-level ISO/IEC 42001 crosswalk, industry-specific legal coverage, or a complete control catalog. Future reviews should assess each required responsibility for an owner, input, decision or activity, output evidence, and feedback path, recording omissions and overlaps. A coverage ratio may then be reported only with a stated requirement set and review procedure; no numeric completeness score is asserted here.

### C Scalability assessment

Scalability was assessed as architectural plausibility: modular APIs, domain-specific agents, delegated ownership, a shared evidence schema, and risk-based escalation could support reuse across business units without central approval of every routine action. The review considered growth in use cases, users, data sources, and cross-system dependencies, identifying potential bottlenecks in approval queues, register synchronization, access enforcement, and monitoring. These are design considerations, not observed capacity or performance results.

Empirical assessment should progressively increase workflow volume, concurrent users, integrations, and participating domains. Measurements should include throughput, percentile response latency, error and recovery rates, evidence freshness, approval waiting time, cost per workflow, and oversight workload. Domain-specific acceptance limits must be defined before testing, with failure and recovery behavior recorded. Multi-unit pilots should also test consistent policy enforcement and delegated accountability. Consequently, the paper claims a potentially extensible architecture while leaving technical and organizational scalability unvalidated.

## III PROPOSED ENTERPRISE AI TRANSFORMATION FRAMEWORK

Figure 2 combines the lifecycle and reference architecture. Business strategy establishes objectives, sponsorship, investment priorities, and a roadmap. Readiness identifies capability gaps before deployment. Governance defines accountability, policies, ethics, compliance, and approval thresholds. Architecture specifies modular APIs, integration, data, security, and infrastructure. Implementation uses staged testing and adoption; operations monitor quality, availability, incidents, and drift. Continuous improvement feeds performance and stakeholder feedback into strategy.




Fig. 2. EATF lifecycle and governed enterprise AI reference architecture.

Eight supporting models operationalize these responsibilities: (1) Enterprise AI Readiness Framework (EARF) assesses organizational, data, process, technology, governance, and workforce readiness; (2) Enterprise AI Governance Model (EAGM) defines policies and lifecycle accountability; (3) Enterprise AI Reference Architecture (EARA) structures experience, agents, workflows, systems, data, and infrastructure; (4) Enterprise AI Workflow Reliability Framework (EAWRF) assesses confidence, accuracy, explainability, human approval, and resilience; (5) Enterprise AI Technical Debt Framework (EATDF) tracks prompt, context, knowledge, model, and automation debt; (6) Enterprise Agent Collaboration Framework (EACF) specifies an orchestrator, domain agents, human supervisor, shared memory, and communication; (7) Enterprise AI Failure Taxonomy (EAF) classifies data, governance, security, architecture, operations, and adoption failures; (8) Enterprise AI Value Realization Framework (EAVRF) measures ROI, productivity, recovery time, customer experience, savings, and adoption.

Data and model documentation support traceability [9], [10]; technical debt motivates maintainability [11]. Agent coordination [12], retrieval from approved knowledge [13], and hallucination controls [14] inform workflow design. Deployment and operational feedback follow established ML engineering practices [15]. ServiceNow, Salesforce, SAP, Microsoft Power Platform, and Oracle are illustrative platforms; EATF prescribes no vendor.

#### IV UNIQUE CONCEPTUAL INNOVATION

The distinctive conceptual contribution proposed in this paper is an evidence-linked enterprise decision architecture: EATF connects business intent, organizational readiness, AI risk controls, technical release evidence, operational feedback, and value realization through one traceable governance structure. Its seven phases and eight supporting models define complementary responsibilities within this structure. The innovation claim concerns this particular synthesis and its decision interfaces; it does not establish that each component is new or that comparable integrations do not exist.

First, EATF proposes a common unit of traceability: an AI-enabled business use case with a persistent identifier linking its owner, intended use, readiness gaps, architecture

dependencies, risk treatments, release versions, approvals, incidents, and benefit measures. This record joins assurance evidence to transformation decisions across workflow platforms. Data and model documentation and auditing practices provide foundations [5], [9], [10]; EATF proposes their connection to adoption and portfolio decisions.

Second, EATF separates and coordinates three decision rights: investment authorization, technical release approval, and residual-risk acceptance. A use case advances only when the applicable business, engineering, and risk conditions are satisfied. Central governance establishes thresholds and escalation rules, while domain teams produce evidence and retain delegated operational authority. This design makes readiness and accountability explicit conditions for deployment rather than treating model performance as sufficient authorization.

Third, EATF closes the feedback loop at the enterprise level. EATF classifies failures, EATDF records remediation obligations, EAWRF reassesses workflow reliability, and EAVRF evaluates whether benefits justify continued investment. Their shared use-case record allows an incident or unmet benefit hypothesis to trigger control revision, workflow redesign, retraining, restricted operation, or retirement. These responses extend the proposed review beyond a model-maintenance loop to the business process and its portfolio rationale.

Together, these mechanisms specify how established disciplines can interact: TOGAF informs architecture, NIST AI RMF informs risk outcomes, ISO/IEC 42001 informs the management system, and MLOps supplies engineering evidence [6], [16]–[18]. EATF contributes a proposed coordination pattern linking those disciplines to readiness, decision authority, and realized value. Its novelty remains a research proposition requiring a broader comparative literature review; its usefulness requires empirical testing. The evaluation in Section VI tests whether the added traceability and decision structure reduce duplicated assurance work and unresolved control exceptions without excessive approval overhead.

#### V COMPARISON WITH EXISTING APPROACHES

Table I distinguishes the proposed integration role of EATF from established architecture, risk-management, management-system, and engineering approaches. TOGAF ADM informs architecture and business alignment [16]; NIST AI RMF 1.0 supplies risk outcomes [17]; ISO/IEC 42001:2023 specifies AI management-system requirements [18]; MLOps supplies engineering practices [6]. COBIT and BPMN additionally inform controls and workflow notation [19], [20]. The mappings below are proposed design interpretations, not validated equivalence or evidence of superiority.

**Table I. Complementary framework scopes**

| Model           | Scope and lifecycle                            | Integration in EATF                      |
|-----------------|------------------------------------------------|------------------------------------------|
| EATF            | Strategy to value                              | Cross-phase governance, adoption, value  |
| TOGAF ADM       | Iterative enterprise architecture              | Architecture and business alignment      |
| NIST AI RMF 1.0 | Govern, Map, Measure, Manage                   | AI risks and trustworthiness             |
| MLOps           | Develop, deploy, monitor, retrain              | Engineering and operational controls     |
| ISO/IEC 42001   | AI management system and continual improvement | Policy to workflow controls and evidence |

### A Relationship to NIST AI RMF

NIST AI RMF 1.0 organizes risk management into Govern, Map, Measure, and Manage. These functions are iterative and context dependent; Govern is cross-cutting rather than a first step completed once [17]. NIST already covers organizational accountability, intended-use context, human oversight, third-party risks, monitoring, and deactivation. EATF therefore does not claim novelty from adding governance or lifecycle coverage. Its proposed distinction is a transformation operating model linking investment prioritization, readiness remediation, workflow architecture, adoption, and business-value review to risk-management outcomes.

The mapping is many-to-many. Govern informs EAGM accountability and decision authority across all seven EATF phases. Map informs strategy, EARF readiness, and EARA architecture by documenting intended use, stakeholders, dependencies, and potential impacts. Measure informs EAWRF testing and operational monitoring. Manage informs implementation gates, incident response, rollback, and continuous improvement. EAVRF adds a portfolio view of costs, adoption, and realized benefits alongside these risk measures; NIST also considers benefits and organizational objectives, so this is an implementation emphasis rather than an exclusive capability.

For example, a favorable risk evaluation would not by itself justify enterprise rollout under the proposed EATF. The sponsor must also confirm readiness, integration ownership, workforce preparation, funding, and a measurable value hypothesis. Conversely, a strong business case cannot override an unacceptable risk assessment. This separates investment authorization from technical release and residual-risk acceptance while connecting their evidence.

### B Operational complement to ISO IEC 42001

ISO/IEC 42001:2023 specifies requirements for establishing, implementing, maintaining, and continually improving an AI management system. Its organizational scope includes responsible AI use and management of risks and opportunities [18]. EATF proposes an implementation pattern within that system: EARF identifies capability gaps; EAGM assigns accountable owners; EARA locates controls in enterprise workflows; EAWRF defines acceptance checks; EATF and EATDF connect failures to corrective work; EAVRF links reviews to benefit realization. These mechanisms illustrate

how requirements can become operating practices, without asserting that ISO lacks operational controls.

A proposed evidence chain links each approved AI use case to its business owner, intended use, applicable policy or requirement, risk assessment, selected control, test evidence, release decision, monitoring record, and corrective action. A shared register uses persistent use-case and release identifiers so engineering teams, risk owners, and management reviewers can inspect the same records. Control implementation is assigned to a named owner, and evidence is refreshed when data, models, integrations, or intended use change. This provides a concrete bridge between management commitments and workflow execution.

For the illustrative service-management deployment, an access policy becomes source-level permissions and restricted agent tools; an oversight requirement becomes approval before consequential ERP or CRM changes; an operational review becomes a review of incidents, control exceptions, adoption, and costs. EATF proposes how to organize these artifacts across platforms. It is not a certification scheme, and using it does not establish ISO/IEC 42001 conformity. A complete requirements and control-applicability assessment remains necessary.

### C Connecting MLOps to enterprise governance

MLOps connects development, validation, deployment, monitoring, and maintenance through engineering roles and automated pipelines [6], [15]. Governance can already be included in mature MLOps implementations. EATF proposes a common enterprise interface for those controls across business units, including purchased AI services and agent workflows that may not train models internally. Central policy defines mandatory evidence and escalation rules; domain teams implement context-specific tests and operations; business and risk owners retain the relevant decision authority.

Three proposed gates connect pipelines to enterprise decisions. At intake, a business owner records purpose, expected value, data authorization, dependencies, and risk tier. At release, engineering supplies versioned data and model documentation, evaluation results, security and access checks, rollback readiness, and workflow tests; the designated business owner and risk authority approve within their delegated limits [9], [10]. During operation, breached thresholds or material changes trigger reassessment, restricted execution, rollback, or suspension. Higher-impact exceptions escalate to an enterprise governance body; routine changes remain with authorized domain owners.

Pipeline evidence is linked to the shared use-case register rather than collected again for each review. Operational telemetry includes model quality and drift, but also failed business transactions, unauthorized tool calls, overrides, user complaints, costs, and workflow recovery. For generative and agentic systems, the release record additionally versions prompts, retrieval sources, tool permissions, and agent configurations, with tests for unsupported outputs and consequential actions [14]. These are proposed EATF extensions of the engineering-to-governance interface, not capabilities attributed to every MLOps practice model.

The resulting decision record joins technical health, risk exposure, workforce adoption, and business outcomes. A technically accurate model may still be paused if its workflow breaches access controls or delivers insufficient value. Retraining or replacing a provider does not automatically authorize production use; the applicable change and release gates must be satisfied. EATF thus specifies who can decide, what evidence is required, and how incidents influence portfolio continuation or retirement.

## VI ILLUSTRATIVE CASE AND EVALUATION

The illustrative case concerns a hypothetical manufacturing enterprise seeking AI-assisted service management across ERP, CRM, and workflow systems. The pilot is bounded to internal service-request classification, retrieval of approved guidance, and drafting of responses or proposed actions. Autonomous approval of financial changes, unrestricted customer-data access, and unapproved cross-system updates are outside the initial pilot boundary. The scenario demonstrates how EATF responsibilities interact; it represents neither an observed organization nor measured results.

The baseline process requires a service analyst to inspect requests, search multiple knowledge sources, consult ERP or CRM owners, and manually record the resolution. Before introducing AI, the proposed pilot would document request volumes and complexity, handling time, rework, escalations, access exceptions, and user satisfaction. This baseline would support a matched-workflow comparison rather than an assumption that automation necessarily improves outcomes.

During readiness assessment, EARF would examine data stewardship, permission consistency, knowledge freshness, API ownership, and staff competency. Missing ownership or unreliable source content would become remediation tasks with an accountable owner and completion evidence. Strategy and EAVRF would define the expected benefit and the cost of integration, assurance, training, and oversight. A sponsor could authorize a limited pilot while deferring enterprise expansion until readiness conditions were satisfied.

In a worked scenario, an employee submits a service request concerning a delayed order. An EACF orchestrator routes it to a classification agent and a retrieval agent operating within EARA's approved data connections. The agents retrieve permitted order status and current service guidance, cite their sources, and draft a response. The service analyst checks the draft and resolves conflicting or missing information. A proposed ERP or CRM modification is routed to the authorized process owner; the tool executes it only after the required approval and records the action against the use-case and release identifiers. This is a proposed workflow, not an executed experiment.

Before release, EAWRF would test incorrect classification, obsolete guidance, unsupported statements, prohibited access, tool failures, and unavailable integrations. Engineering would attach evaluation and rollback evidence; EAGM would identify who can approve the release and accept residual risk. Low-confidence or unsupported requests would be routed to manual handling under predefined rules. These controls illustrate the distinction between an agent producing a

suggestion and an accountable owner authorizing a business action.

If monitoring later detected an unauthorized tool attempt or an incorrect action, the affected capability would be restricted or paused, relevant records preserved, and the incident routed to its owner. EAFT would classify the failure; EATDF would track remediation to the prompt, retrieval source, permissions, model, or integration; and EAWRF would require regression checks before restoration. EAVRF would review whether benefit, user adoption, and assurance costs support continued use. The plausible outcome is a traceable decision process; faster handling, fewer errors, or lower costs remain hypotheses to test.

To test the incremental contribution, a future pilot should compare an existing NIST/ISO-informed MLOps baseline with that baseline plus the EATF register, decision gates, and portfolio review. Measures should include evidence completeness, duplicate assurance effort, time to accountable approval, unresolved control exceptions, workflow incidents, adoption, and benefit realization. Comparable workflows and observation periods, implementation costs, and organizational differences should be documented. A null or adverse result must be reported; conceptual integration alone does not demonstrate improved outcomes.

Future pilots should compare matched workflows and account for complexity and adoption. Measures include strategic alignment and sponsorship; data quality, integration readiness, workforce competency, and process maturity; policy compliance and audit findings; architecture conformance, API reuse, and scalability; deployment success, defects, time to production, and active use; availability, accuracy, incidents, and mean time to recovery; ROI, savings, productivity, customer satisfaction, and innovation. Readiness composites require validated weights. Planning/reporting effort, decision latency, dependency visibility, and risk lead time can test workflow benefits. No percentage improvements are asserted.

## VII DISCUSSION AND CONCLUSION

EATF connects organizational preparation, delivery, assurance, and value in a modular research and practitioner roadmap. Proposed benefits include coherent governance, knowledge reuse, and reduced fragmentation. Risks include poor data, bias, security failures, weak sponsorship, resistance, legacy constraints, lock-in, insufficient testing, and drift. Mitigations combine stewardship, training, modular architecture, staged release, accountable review, and monitoring. Conceptual status, sector variation, and absent empirical evidence limit generalization. Future work should validate maturity scores, conduct longitudinal multi-sector comparisons, and refine agent governance and automated assessments. EATF's proposed contribution is the evidence-linked coordination of transformation, assurance, and portfolio decisions through shared use-case records and explicit decision rights. This conceptual innovation remains subject to comparative review and empirical validation; no new algorithm or measured performance gain is claimed.

## REFERENCES

- [1] G. Vial, "Understanding digital transformation: A review and a research agenda," *J. Strategic Information Systems*, vol. 28, no. 2, pp. 118–144, 2019. doi: 10.1016/j.jsis.2019.01.003.
- [2] J. Jöhnk, M. Weißert, and K. Wyrki, "Ready or Not, AI Comes—An Interview Study of Organizational AI Readiness Factors," *Business & Information Systems Engineering*, vol. 63, pp. 5–20, 2021. doi: 10.1007/s12599-020-00676-7.
- [3] P. Mikalef and M. Gupta, "Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance," *Information & Management*, vol. 58, no. 3, art. 103434, 2021. doi: 10.1016/j.im.2021.103434.
- [4] A. Jobin, M. Ienca, and E. Vayena, "The global landscape of AI ethics guidelines," *Nature Machine Intelligence*, vol. 1, pp. 389–399, 2019. doi: 10.1038/s42256-019-0088-2.
- [5] I. D. Raji et al., "Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing," *Proc. FAT\**, pp. 33–44, 2020. doi: 10.1145/3351095.3372873.
- [6] D. Kreuzberger, N. Kühn, and S. Hirschl, "Machine Learning Operations (MLOps): Overview, Definition, and Architecture," *IEEE Access*, vol. 11, pp. 31866–31879, 2023. doi: 10.1109/ACCESS.2023.3262138.
- [7] A. R. Hevner, S. T. March, J. Park, and S. Ram, "Design Science in Information Systems Research," *MIS Quarterly*, vol. 28, no. 1, pp. 75–105, 2004. doi: 10.2307/25148625.
- [8] K. Peffers, T. Tuunanen, M. A. Rothenberger, and S. Chatterjee, "A Design Science Research Methodology for Information Systems Research," *J. Management Information Systems*, vol. 24, no. 3, pp. 45–77, 2007. doi: 10.2753/MIS0742-1222240302.
- [9] T. Gebru et al., "Datasheets for datasets," *Communications of the ACM*, vol. 64, no. 12, pp. 86–92, 2021. doi: 10.1145/3458723.
- [10] M. Mitchell et al., "Model Cards for Model Reporting," *Proc. FAT\**, pp. 220–229, 2019. doi: 10.1145/3287560.3287596.
- [11] D. Sculley et al., "Hidden Technical Debt in Machine Learning Systems," *Advances in Neural Information Processing Systems*, vol. 28, pp. 2503–2511, 2015.
- [12] M. Wooldridge and N. R. Jennings, "Intelligent agents: Theory and practice," *Knowledge Engineering Review*, vol. 10, no. 2, pp. 115–152, 1995. doi: 10.1017/S0269888900008122.
- [13] P. Lewis et al., "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," *Advances in Neural Information Processing Systems*, vol. 33, pp. 9459–9474, 2020.
- [14] Z. Ji et al., "Survey of Hallucination in Natural Language Generation," *ACM Computing Surveys*, vol. 55, no. 12, art. 248, 2023. doi: 10.1145/3571730.
- [15] S. Amershi et al., "Software Engineering for Machine Learning: A Case Study," *Proc. ICSE-SEIP*, pp. 291–300, 2019. doi: 10.1109/ICSE-SEIP.2019.00042.
- [16] The Open Group, *The TOGAF Standard*, 10th ed., 2022. Available: [opengroup.org/togaf](https://opengroup.org/togaf).
- [17] NIST, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, 2023. doi: 10.6028/NIST.AI.100-1. Available: <https://airc.nist.gov/airmf-resources/airmf/5-sec-core/>. Version compared: 1.0.
- [18] ISO/IEC, *Information technology—Artificial intelligence—Management system*, ISO/IEC 42001:2023, 2023. Overview: <https://www.iso.org/standard/42001>.
- [19] ISACA, *COBIT 2019 Framework: Introduction and Methodology*, 2018.
- [20] Object Management Group, *Business Process Model and Notation*, ver. 2.0.2, 2014. Available: [omg.org/spec/BPMN/2.0.2](https://omg.org/spec/BPMN/2.0.2).