INTELLIGENT SYSTEMS PRESS™
EXECUTIVE ASSESSMENT SERIES

EA-011: AI Vendor & Third-Party Risk Assessment™

How Effectively Does Your Enterprise Govern AI Risk Beyond Its Direct Control?

Purpose

A practical executive assessment for evaluating how effectively the enterprise selects, contracts with, integrates, monitors, and exits AI vendors and other third parties whose technologies, data, models, services, or personnel can create material operational, security, privacy, compliance, ethical, financial, and reputational exposure.

Educational Executive Self-Assessment | Version 1.0

www.intelligentsystemspress.com

About This Assessment

Enterprises increasingly depend on external AI platforms, embedded features, cloud services, data providers, model providers, consultants, and subcontractors. Accountability, however, cannot be outsourced. This assessment helps leaders determine whether third-party AI dependencies are governed with sufficient due diligence, contractual clarity, technical validation, operational oversight, transparency, resilience, and exit readiness.

How to Complete It

1. Rate each statement	Use the 1-5 scale based on current, observable enterprise practice - not procurement intent, contract language alone, vendor claims, or isolated reviews.
2. Score each domain	Add the five ratings in each domain. Every domain score ranges from 5 to 25.
3. Calculate the total	Add all eight domain scores. The overall AI vendor and third-party risk score ranges from 40 to 200.
4. Interpret the pattern	Review the total and the distribution. Any domain below 15 represents a material dependency that may expose the enterprise to unmanaged third-party risk.
5. Select 90-day actions	Choose a limited number of executive actions that strengthen inventory, due diligence, contracts, integration controls, monitoring, incident coordination, resilience, or exit readiness.

Rating Scale

1 Strongly Disagree	Practice is absent or contradicted by current evidence.
2 Disagree	Practice is inconsistent, informal, or limited to isolated areas.
3 Partly Agree	Practice exists in several areas but is not yet enterprise-wide.
4 Agree	Practice is generally reliable, measured, and consistently applied.
5 Strongly Agree	Practice is integrated, evidence-driven, and continuously improved.

Recommended use: Complete individually first, then compare results across business, procurement, legal, privacy, cybersecurity, technology, risk, compliance, operations, and internal audit leaders. Differences in ratings are evidence of accountability or control misalignment and should be discussed before averaging scores.

Domain 1: Third-Party AI Inventory & Accountability

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree). Base the rating on repeatable enterprise behavior and current evidence.

The enterprise maintains a current inventory of external AI products, embedded AI features, model providers, data providers, consultants, subcontractors, and material downstream dependencies.					
Each third-party AI relationship has a named business owner and accountable risk owner with authority to approve, restrict, suspend, or terminate use.					
The inventory identifies intended use, affected processes, data exchanged, decisions influenced, users, jurisdictions, hosting arrangements, criticality, and risk classification.					
Unapproved or informally acquired AI services can be detected, investigated, and brought under governance or removed in a timely manner.					
Changes in vendor ownership, service scope, model architecture, subprocessors, data practices, or embedded AI functionality trigger inventory and risk updates.					

Domain score: _____ / 25

Evidence or observations:

Domain 2: Due Diligence & Risk Classification

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree). Base the rating on repeatable enterprise behavior and current evidence.

Risk-based due diligence is completed before selection and is proportionate to the AI use case, data sensitivity, decision impact, operational criticality, and stakeholder exposure.					
Vendor claims about performance, security, privacy, fairness, explainability, compliance, resilience, and responsible AI are validated with evidence appropriate to the risk.					
Due diligence examines model and data provenance, development practices, testing, limitations, known incidents, regulatory exposure, financial stability, and reliance on subcontractors.					
High-risk or opaque vendor offerings receive enhanced legal, technical, cybersecurity, privacy, compliance, ethics, and operational review before approval.					
Procurement decisions document accepted risks, required controls, unresolved concerns, compensating measures, and the executive authority responsible for approval.					

Domain score: _____ / 25

Evidence or observations:

Domain 3: Contractual Protections & Rights

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree). Base the rating on repeatable enterprise behavior and current evidence.

Contracts clearly define approved uses, prohibited uses, service boundaries, data ownership, confidentiality, intellectual-property rights, and accountability for AI-enabled outcomes.					
Agreements establish requirements for security, privacy, retention, deletion, model training, data reuse, subprocessors, jurisdiction, and regulatory cooperation.					
The enterprise has enforceable rights to receive material information about model changes, limitations, incidents, control failures, audits, and changes in subprocessors or data practices.					
Service levels, performance expectations, human-oversight requirements, remediation obligations, indemnification, liability allocation, and termination rights reflect the risk of the use case.					
Contracts preserve sufficient audit, testing, access, portability, continuity, and exit rights to prevent the enterprise from depending solely on vendor assurances.					

Domain score: _____ / 25

Evidence or observations:

Domain 4: Technical Integration & Data Controls

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree). Base the rating on repeatable enterprise behavior and current evidence.

Third-party AI services are integrated through approved architectures that enforce identity, access, least privilege, encryption, logging, segmentation, and secure configuration.					
Data shared with vendors is minimized, authorized, classified, protected, and limited to the approved purpose and retention period.					
Inputs, outputs, interfaces, prompts, plugins, agents, APIs, and automated actions are tested for misuse, leakage, injection, unauthorized access, and unintended downstream effects.					
Enterprise controls can restrict or disable vendor features, model versions, data connections, automated actions, and user access when risk exceeds approved thresholds.					
Changes to vendor models, interfaces, configurations, integrations, or data flows are assessed and validated before material production impact occurs.					

Domain score: _____ / 25

Evidence or observations:

Domain 5: Performance, Safety & Responsible-Use Validation

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree). Base the rating on repeatable enterprise behavior and current evidence.

Vendor AI performance is independently evaluated using enterprise-relevant data, scenarios, user groups, operating conditions, and decision thresholds before production use.					
Testing addresses accuracy, reliability, robustness, harmful output, bias, explainability, accessibility, human factors, and foreseeable misuse in proportion to risk.					
Known limitations, confidence boundaries, failure modes, and conditions requiring human review are documented and communicated to affected users and decision-makers.					
Acceptance criteria prevent a vendor solution from entering or remaining in production when evidence does not support its intended purpose or risk tolerance.					
Revalidation occurs after material model changes, data shifts, new use cases, incidents, performance degradation, or changes in legal and operating conditions.					

Domain score: _____ / 25

Evidence or observations:

Domain 6: Ongoing Monitoring & Assurance

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree). Base the rating on repeatable enterprise behavior and current evidence.

The enterprise continuously monitors vendor service availability, performance, security, privacy, compliance, responsible-use indicators, and operational impact at a frequency matched to risk.					
Vendor attestations, certifications, audit reports, testing evidence, control reports, and remediation commitments are reviewed critically rather than accepted as substitutes for oversight.					
Material changes in models, policies, ownership, subprocessors, data use, geography, pricing, support, or control posture trigger reassessment and approval when required.					
Business owners and oversight functions receive actionable reporting on vendor risk, exceptions, incidents, unresolved findings, concentration exposure, and remediation status.					
Independent assurance or targeted testing is used for high-risk relationships when internal evidence or vendor transparency is insufficient.					

Domain score: _____ / 25

Evidence or observations:

Domain 7: Incident Coordination & Resilience

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree). Base the rating on repeatable enterprise behavior and current evidence.

Contracts and operating procedures define how the vendor and enterprise will detect, report, investigate, contain, communicate, remediate, and learn from AI-related incidents.					
Escalation paths, notification timeframes, decision authority, evidence preservation, regulatory coordination, customer communication, and public-response responsibilities are established in advance.					
Business continuity plans address vendor outages, degraded performance, corrupted outputs, compromised data, unavailable models, loss of support, and abrupt service changes.					
The enterprise tests contingency procedures, manual alternatives, fail-safe behavior, data recovery, and coordinated incident response for material third-party AI dependencies.					
Lessons from incidents, near misses, vendor failures, and industry events are incorporated into controls, contracts, monitoring, training, and sourcing decisions.					

Domain score: _____ / 25

Evidence or observations:

Domain 8: Concentration, Portability & Exit Readiness

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree). Base the rating on repeatable enterprise behavior and current evidence.

Leadership understands concentration risk created by dependence on a single provider, model family, cloud platform, data source, integration pattern, or specialized vendor capability.					
Critical AI services have proportionate alternatives, fallback procedures, internal capabilities, or contingency arrangements that reduce disruption from vendor failure or strategic change.					
Data, configurations, prompts, workflows, logs, documentation, and other necessary assets can be retrieved or transferred in usable formats when the relationship ends.					
Exit plans address transition responsibilities, service continuity, access revocation, data return or destruction, model dependencies, records retention, user communication, and control validation.					
Vendor relationships are periodically renewed, renegotiated, restricted, replaced, or retired based on current value, risk, performance, transparency, resilience, and strategic fit.					

Domain score: _____ / 25

Evidence or observations:

Scoring & Third-Party Risk Interpretation

Record each domain score. Treat any domain below 15 as a priority dependency even when the overall score is higher.

1. Third-Party AI Inventory & Accountability	____ / 25
2. Due Diligence & Risk Classification	____ / 25
3. Contractual Protections & Rights	____ / 25
4. Technical Integration & Data Controls	____ / 25
5. Performance, Safety & Responsible-Use Validation	____ / 25
6. Ongoing Monitoring & Assurance	____ / 25
7. Incident Coordination & Resilience	____ / 25
8. Concentration, Portability & Exit Readiness	____ / 25
Total AI vendor and third-party risk score	**____ / 200**

Overall Third-Party Risk Capability Levels

180-200	**Third-Party Assured**	Vendor governance, evidence, contracts, integration controls, monitoring, resilience, and exit readiness operate as an integrated and continuously improved system.
150-179	**Risk Aligned**	Strong third-party controls exist, with targeted gaps in transparency, contract rights, validation, monitoring, concentration, or contingency planning.
110-149	**Developing**	Important practices are present, but due diligence, contractual protection, technical validation, oversight, resilience, or exit readiness remains uneven.
70-109	**Fragmented**	Third-party AI risk depends on isolated reviews or vendor assurances, with material weaknesses in ownership, evidence, controls, monitoring, or continuity.
40-69	**Exposed**	Foundational inventory, accountability, due diligence, contractual, technical, assurance, incident, and exit mechanisms are insufficient for responsible reliance on third parties.

Important: This assessment is a discussion and prioritization tool, not a certification or substitute for vendor due diligence. Results should be validated using current inventories, contracts, architecture and data-flow records, security and privacy reviews, testing evidence, audit reports, incident records, service performance, vendor communications, and exit documentation.

Executive Reflection & 90-Day Action Plan

Use the scoring pattern to identify the few leadership actions that will most reduce unmanaged dependence on external AI providers and strengthen accountability across the full relationship lifecycle.

Third-party strengths

Which inventory, due-diligence, contract, integration, validation, monitoring, resilience, or exit practices currently provide the strongest assurance?

Critical exposures

Which vendors, use cases, data flows, contractual gaps, opaque dependencies, concentration risks, or continuity weaknesses create the greatest enterprise exposure?

Decisions required from leadership

What sourcing, contractual, investment, risk-acceptance, remediation, suspension, diversification, continuity, or exit decisions cannot be delegated?

90-Day Priority Actions

Executive commitment: We will review progress by _____ and reassess affected domains using current third-party risk evidence.

Continue Your Third-Party AI Governance Journey

Responsible AI extends beyond the enterprise boundary. Sustainable value requires clear ownership, evidence-based selection, enforceable rights, secure integration, independent validation, ongoing oversight, coordinated resilience, and credible exit options.

Explore the Intelligent Systems Press collection and Free Resources library
www.intelligentsystemspress.com