Journal of Data Processing

DLINE JOURNALS

Print ISSN: 2278 – 6481
Online ISSN: 2278 – 649X

JDP 2026: 16 (4)
https://doi.org/10.6025/jdp/2026/16/4/231-267

The Architecture of Trust: Mapping GRC AI Agents, Verified Auditors, and the Governance Gap in Enterprise AI Ecosystems

Hathairat Ketmaneechairat
Department of Information and Production Technology Management
College of Industrial Technology
King Mongkut's University of Technology North Bangkok
Thailand
hathairatk@kmutnb.ac.th

ABSTRACT

Artificial intelligence (AI) is transforming business to business (B2B) industrial marketing, procurement, and governance, shifting operations from static analytics to autonomous, agentic workflows. However, because AI functions as an actor that rearranges inter organizational interfaces, its rapid deployment introduces significant relational and governance risks. This study investigates the tension between agentic AI capabilities and the "accountability assets" required for enterprise trust by mapping the dual ecosystems of Governance, Risk, and Compliance (GRC) AI agents and verified compliance auditors. Analyzing a curated directory of 407 AI-agent and 503 audit firm listings, we employ descriptive statistics, cross-ecosystem framework mapping, Natural Language Processing (NLP) topic modeling, and exploratory clustering. Our findings reveal a profound "verification gap" and ecosystem asymmetry: while 100% of verified audit firms list compliance frameworks, only 20.1% of AI agents possess structured, verifiable framework records. AI agents exhibit high functional siloing (e.g., application security, AML, SecOps), with "compliance" agents acting as a critical bridge to governance verification. Geographically, a distinct division of labor emerges, with AI innovation concentrated in the United States and the labor-intensive audit ecosystem heavily distributed across global talent pools such as India. We conclude that enterprise AI adoption is fundamentally a governance challenge. B2B buyers must look beyond model capabilities to evaluate vendors' accountability assets, favoring agents with verifiable alignments with independent auditor frameworks (e.g., ISO 27001, SOC 2, FedRAMP) and architecting multi-agent pipelines with layered runtime guardrails.

Keywords: Artificial Intelligence, Agentic AI, Governance, Risk, and Compliance (GRC), B2B Industrial Marketing, Enterprise AI Architecture, Accountability Assets, Compliance Auditing, Algorithmic Governance

Received: 4 May 2026, Revised: 12 July 2026, Accepted 7 August 2026

Copyright: DLINE

1. Introduction

Artificial intelligence (AI) is transforming industrial marketing at both the exchange level and the firm capability level. Yet the body of empirical evidence in business to business (B2B) settings remains uneven and frequently lacks robust theoretical underpinning. Research published in *Industrial Marketing Management* has drawn attention to the relational and governance risks that arise when AI enters inter organizational interfaces. AI is not merely a tool; it functions as an actor that rearranges relationships among buying and selling centers through mechanisms such as algorithmic gatekeeping, reduced transparency, and shifts in power [1, 2].

These relational insights are complemented by capital market evidence. The first large-sample event study of B2B marketing finds favorable yet heterogeneous shareholder responses to AI-enabled initiatives. The magnitude of value creation is conditional on environmental turbulence and customer base complexity, thereby delineating clear boundary conditions for the returns to AI in industrial markets [3]. At the organizational buying level, conceptual and editorial work documents a parallel shift toward buying ecosystems. This shift raises new questions about multi-actor information flows and the governance of AI-augmented procurement and solution selling processes [4, 5]. More broadly, state of the art reviews in the marketing literature trace how generative artificial intelligence (GenAI) is reshaping innovation and go to market practices. They call for research that links firm level decisions to adopt GenAI with subsequent value appropriation and stakeholder reactions an integrated research program that industrial marketing has only begun to pursue [6, 7, 8].

Several theoretical orientations offer relevant foundations, although they remain only partially integrated. At the individual user level, adoption is commonly framed through the Unified Theory of Acceptance and Use of Technology (UTAUT), whose core constructs performance expectancy, effort expectancy, social influence, and facilitating conditions capture key drivers of uptake [9]. The Technology Acceptance Model (TAM) frequently serves as a complementary item source, with performance expectancy and effort expectancy mapping directly onto the corresponding UTAUT dimensions [10, 11]. At the firm level, the Technology–Organization Environment (TOE) framework explains how technological readiness, organizational support, and regulatory pressure jointly shape both the decision to adopt and the depth of subsequent implementation [12, 13, 14].

AI is especially attractive in environments that generate large volumes and diverse streams of data from development tools commit histories, pull-request activity, build logs, access records, and runtime telemetry [15]. Nevertheless, many current AI implementations remain isolated, reactive, and dependent on manual interpretation. Agentic practices supported by AI offer a promising alternative [16]. Unlike fixed analytics pipelines, agentic systems enable semi-autonomous or autonomous software agents to detect signals, reason about context, make decisions, and act across interconnected systems. In an insider threat scenario, for example, an agentic security layer can continuously observe behaviors, correlate cross repository and pipeline signals, compute risk in real time, enact adaptive access decisions, and trigger responses such as session termination, secrets rotation, privileged access revalidation, or incident-ticket creation [17, 18]. Organizations can further strengthen continuous mitigation by coordinating specialized agents user behavior analytics agents, anomaly detection agents, data loss prevention agents, NLP-based communication monitoring agents (where policy permits), and incident response agents [19].

Yet AI and agent based solutions introduce significant challenges: limited transparency and explainability, elevated false positive rates and alert fatigue, privacy and ethical constraints, potential bias in behavioral models, vulnerability to adversarial manipulation, and the consequent need for robust governance [20]. Many AI-enabled capabilities also suffer from weak appropriability. General purpose models can readily approximate classification, extraction, summarization, drafting, and first pass reasoning [21]. Where appropriability is weak, control over cospecialized complementary assets becomes the principal route to durable value capture [22]. Although the classic framework was developed primarily for tangible commercialization assets that convert inventions into usable products, AI-supported work introduces a distinct form of cospecialized asset. Accountability assets do not commercialize an innovation; they render its outputs

usable and defensible in organizational settings in which someone must ultimately stand behind the decision [23].

Although AI agents are increasingly positioned to automate governance, risk, and compliance activities, their functional capabilities and their demonstrable compliance status represent distinct dimensions of enterprise readiness. Similarly, the existence of independent audit and assessment organizations does not, by itself, establish that the capabilities offered by AI providers are supported by appropriate and verifiable assurance mechanisms. Understanding this distinction requires a framework that connects AI-agent functionality, documented compliance framework coverage, and independent assurance without treating them as interchangeable constructs. The following section develops this conceptual foundation and identifies the empirical questions that can be addressed through the available provider directory data.

2. Conceptual Framework: Agent Capabilities and Assurance Requirements

2.1 Integrated Evaluation Implications

The enterprise AI architecture is best understood as a coordinated system in which data preparation, domain adaptation, knowledge retrieval, safety controls, and deployment contribute different but interdependent functions. The distinction between model capability and system level reliability is central to this perspective. Domain adaptation may improve task alignment without ensuring that the model has access to current information; retrieval-augmented generation (RAG) may provide external evidence without guaranteeing that the retrieved material is relevant or correctly interpreted; and guardrails may identify or constrain selected forms of problematic behavior without eliminating every potential failure. Similarly, scalable inference and application integration facilitate operational use but do not, by themselves, establish factual accuracy, privacy protection, security, or regulatory compliance. The overall suitability of an enterprise AI system therefore depends on the interaction between its functional capabilities and the controls and evidence available to support their responsible use.

This integrated perspective informs the empirical study by distinguishing documented provider capabilities from the assurance evidence needed to evaluate their suitability for organizational use. At the conceptual level, data related functions imply requirements concerning data quality and privacy; task oriented capabilities require task relevant performance evidence; retrieval functions require attention to source relevance and evidential support; safety mechanisms require evaluation against defined risks; and deployment functions require appropriate operational monitoring and access controls. These are evaluation requirements derived from the conceptual framework, not measurements of the systems represented in the directories.

The present study examines the observable information available in the AI-agent and audit firm datasets rather than conducting technical tests of individual systems. Accordingly, the empirical analysis focuses on provider characteristics, recorded tasks and capabilities, framework references, accreditation related information, integration records, geographic distributions, and the completeness and provenance of directory information. Descriptive and comparative analyses establish the characteristics of the two ecosystems; framework mapping examines recorded coverage and overlap; statistical tests assess associations between selected categorical characteristics where the data support such comparisons; and text based analyses explore provider descriptions and potential capability profiles. Missingness and differences in record structure are considered when interpreting these results.

This distinction also defines the limits of the framework's empirical application. The presence of a framework reference does not establish that a provider has implemented every associated control or passed an independent assessment. Likewise, an empty field does not establish that the corresponding capability, integration, or assurance mechanism is absent. Because the study uses directory records, its findings concern documented ecosystem characteristics and observable alignment rather than independently measured technical effectiveness, verified interoperability, organizational adoption, or actual trustworthiness. Demonstrating those outcomes would require additional evidence, such as system level benchmarks, implementation records, or independent assessments of specific AI agents and deployments.

3. Data and Methods

3.1 Data Sources and Study Scope

This study examines two related but distinct components of the enterprise AI governance ecosystem: AI agents and platforms relevant to governance, risk, and compliance (GRC), and independent audit and assessment firms with recorded compliance related frameworks or accreditation information. The analysis uses the Agent Index dataset, [24] described in the source material as an open, continuously maintained, source-verified directory published under the Creative Commons Attribution 4.0 International (CC BY 4.0) licence by Agent Index (Heritage Services).

The supplied records comprise 407 AI-agent listings and 503 audit-firm listings. The two datasets are analysed as distinct provider populations because their entities, attributes, and reporting structures differ. The agent records primarily describe provider identity, categories, tasks, pricing, framework references, integrations, and geographic information. The audit firm records describe provider identity, headquarters and office locations, frameworks, accreditation related information, partner tools, and selected institutional attributes. The comparison therefore concerns the characteristics recorded for each ecosystem, rather than a one to one comparison of equivalent organizations or capabilities.

The dataset is a curated directory, not a demonstrated census of all AI agents or audit firms worldwide. Consequently, the findings describe the providers represented in the supplied records. They should not automatically be interpreted as estimates of total market size, market share, adoption, or the prevalence of capabilities among all providers. The directory's source verification status also needs to be distinguished from independent verification of every capability, framework implementation, or assurance relationship described in an individual record.

3.2 Data Structure and Analytical Units

The original files were retained, and separate normalized analysis tables were used to represent the different entities and relationships in the data. The principal unit of analysis is one AI-agent listing or one audit firm listing. Additional relational tables represent provider framework associations, agent task associations, firm accreditation records, and provider source associations. This structure accommodates the fact that a single provider may be associated with multiple tasks, frameworks, integrations, offices, or accreditation records.

For AI agents, the principal variables include provider identifiers and names, headquarters country and region, primary and secondary categories, pricing model, recorded tasks, framework references, claimed frameworks, integrations, authorizations, India-related relevance, and source check dates. For audit firms, the principal variables include provider identifiers and names, headquarters country, listed offices, framework references, accreditation related records, partner tools, government body information, and source check dates. Framework records retain both the provider association and the normalized framework identifier, while source records preserve the available source URLs and verification dates.

The separation of provider level records from one to many associations prevents multiple tasks, frameworks, offices, or accreditation records from being treated as if they represented separate providers. It also enables framework prevalence and provider characteristics to be calculated using an explicitly defined unit of analysis. Detailed field definitions, original variable names, coding categories, and table schemas are provided in the supplementary data dictionary.

3.3 Data Cleaning, Normalization, and Provenance

Data preparation was designed to improve comparability between the two datasets while preserving the original records for auditability. Framework labels were standardized to address differences in capitalization, spacing, or notation for example, variants such as "ISO 27001" and "iso27001" while retaining the original labels. Country names were mapped to common identifiers so that alternative representations, such as "United States" and "US," could be compared consistently. These transformations support descriptive and comparative analyses without replacing the source values.

Framework fields were treated according to their recorded meaning. In particular, the agent fields labelled frameworks and frameworks_claimed were retained separately rather than automatically combined. A recorded or standardized framework reference and a vendor-claimed framework reference are not necessarily equivalent forms of evidence. Their interpretation depends on the definitions and coding rules associated with the source fields; therefore, normalization of a framework name does not itself establish implementation, certification, or independent verification.

Potential duplicate records were examined using identifiers such as slugs, provider names, vendor identities, and source URLs. Similar names alone were not treated as sufficient grounds for merging records. Source URLs, available source-check dates, original values, and normalization decisions were retained to preserve provenance and support review of the resulting analysis tables. The detailed normalization mappings and field level cleaning decisions are documented in the supplementary material.

3.4 Missingness and Evidence Interpretation

Missing or empty fields constitute an important feature of the supplied records and were examined before interpreting framework coverage or integration patterns. In the initial data inspection, 325 of the 407 AI-agent records had empty frameworks lists, 270 had empty frameworks_claimed lists, and 391 had no recorded integrations. Among the 503 audit firm records, 478 had no recorded partner_tools. These counts demonstrate substantial sparsity in selected directory fields and indicate that observed coverage depends on the field examined and the denominator used.

An empty list was therefore interpreted as information not recorded in the directory, rather than automatically as a verified absence of the corresponding capability or relationship. Similarly, a non empty field was treated as a recorded attribute whose evidential status depends on its source and definition. This distinction is particularly important when comparing framework references across the two ecosystems, because differences in directory structure or reporting practices may contribute to apparent differences in coverage.

Framework prevalence was consequently interpreted with explicit attention to the denominator. Analyses using all provider records describe the prevalence of recorded framework information in the full directory, whereas analyses restricted to providers with at least one framework reference describe the distribution conditional on framework information being present. These quantities answer different questions and should not be conflated. The study therefore considers the full-record denominator alongside conditional descriptions where relevant, with the corresponding interpretation stated explicitly.

The missingness assessment also limits the strength of conclusions that can be drawn from cross-ecosystem comparisons. A higher proportion of recorded frameworks in one dataset does not, by itself, demonstrate superior assurance capacity, and an unrecorded integration does not prove that integration is unavailable. Findings are accordingly reported as patterns of documented coverage and information completeness. The supplementary material provides the detailed field level definitions, normalization mappings, and sensitivity scenarios used to support these interpretations.

3.5 Analytical Strategy

The analytical strategy follows the study's research questions and proceeds from dataset description to cross-ecosystem comparison. Descriptive statistics characterize the provider populations, categories, tasks, pricing models, framework references, geographic distributions, and selected record-level counts. Comparative analyses examine framework prevalence, framework overlap, and co-occurrence patterns while preserving the distinction between the agent and audit-firm units of analysis. Association tests are applied to selected categorical or count variables where the sample structure and test assumptions permit, with effect sizes reported alongside significance measures where appropriate.

Geographic analyses summarize headquarters and office distributions and examine the representation of India in the supplied records. Text based analyses use provider descriptions to explore recurring themes and potential provider profiles. Any clustering or topic-based results are interpreted according to the validation

actually performed; internal separation metrics alone are not treated as evidence of semantic validity or external generalizability. Data completeness, normalization choices, and differences in the treatment of recorded and unrecorded information are considered when interpreting the results.

All analyses are descriptive of the supplied directory records unless an inferential procedure is explicitly identified. Statistical associations do not establish causality, and directory level patterns do not independently demonstrate operational effectiveness, real world adoption, or the quality of a specific assurance engagement. The following subsections specify the individual analytical procedures, variables, statistical tests, and outputs used to address each research question.

4. Results

Dataset	Field	Meani	N	Missing/ empty n	Missing/ empty %	empty n
AI agents	slug	Stable record identifier	407	0	0	407
AI agents	vendor	Vendor/provider	407	0	0	407
AI agents	hq_country	Headquarters country	407	0	0	407
AI agents	hq_region	Headquarters region	407	0	0	407
AI agents	category	Primary category	407	0	0	407
AI agents	categories	All listed categories	407	0	0	407
AI agents	pricing	Pricing text	407	31	7.6	376
AI agents	pricing_model	Pricing model class	407	0	0	407
AI agents	frameworks	Normalized/ structured frameworks	407	325	79.9	82
AI agents	frameworks_ claimed	Frameworks claimed in source wording	407	270	66.3	137
AI agents	tasks	Recorded tasks	407	0	0	407

AI agents	india_relevance	India-related relevance note	407	331	81.3	76
AI agents	authorizations	Authorization records	407	335	82.3	72
AI agents	integrations	Recorded integrations	407	391	96.1	16
AI agents	summary	Summary text	407	0	0	407
AI agents	description	Description text	407	0	0	407
AI agents	sources	Source references	407	0	0	407
AI agents	last_checked	Source check date	407	0	0	407
Audit firms	slug	Stable record identifier	503	0	0	503
Audit firms	hq_country	Headquarters country	503	0	0	503
Audit firms	offices	Office locations	503	0	0	503
Audit firms	frameworks	Listed frameworks	503	0	0	503
Audit firms	accreditations	Accreditation/ source evidence records	503	0	0	503
Audit firms	partner_tools	Recorded partner tools	503	0	0	503
Audit firms	summary	Summary text	503	478	95	25
Audit firms	sources	Source references	503	0	0	503
Audit firms	last_checked	Source check date	503	0	0	503
Audit firms	government_body	Government-body indicator	503	0	0	503

Table 1. Dataset characteristics, variables and missingness

Missing/empty means null, empty string or empty list; boolean False is not treated as missing

This table summarises the structure, completeness and variable definitions of the two source datasets (AI agents, n = 407; audit firms, n = 503). For every recorded field, it reports the total number of records (N), the count and percentage of missing or empty values (null, empty string or empty list), and the corresponding non-empty count. Boolean false values are not treated as missing. The table demonstrates near-complete coverage for core identifiers (slug, vendor/headquarters, category, tasks, summary, sources, last_checked) while highlighting substantial missingness in optional or sparsely populated fields such as structured frameworks (79.9 % empty among agents), integrations (96.1 % empty) and partner tools (95 % empty among auditors).

Dataset	Records	Unique slug	Unique HQ Coun-tries	Me-dian frame works/ record	Me-dian frame works / record	Me-dian sources / record	Me-dian sources / record	Me-dian text words	Records with frame works	Records with frame works %
AI agents	407	407	30	0	0.48	4	4.36	36	82	20.1
Audit firms	503	503	15	1	1.51	3	3.41	16	503	100

Table 2. Descriptive statistics for AI agents and audit firms

Aggregate descriptive statistics are presented for both datasets. Metrics include total records, unique slugs, unique headquarters countries, median and mean number of frameworks per record, median and mean number of sources per record, median text word count, and the absolute and percentage counts of records that contain at least one framework entry. Agents show low framework density (median 0, mean 0.48; only 20.1 % of records list any structured framework), whereas every audit-firm record lists at least one framework (median 1, mean 1.51).

Dataset	Dimension	Label	Count	Percent of records (%)
AI agents	Primary category	security	117	28.7
AI agents	Primary category	business-ops	109	26.8
AI agents	Primary category	risk	89	21.9
AI agents	Primary category	compliance	53	13
AI agents	Primary category	code-security	39	9.6
AI agents	Pricing model	quote	321	78.9
AI agents	Pricing model	published-tiers	62	15.2

AI agents	Pricing model	freemium	20	4.9
AI agents	Pricing model	annual-contract	4	1
AI agents	HQ country	United States	173	42.5
AI agents	HQ country	United Kingdom	41	10.1
AI agents	HQ country	Germany	34	8.4
AI agents	HQ country	India	24	5.9
AI agents	HQ country	France	22	5.4
AI agents	HQ country	Israel	16	3.9
AI agents	HQ country	Japan	11	2.7
AI agents	HQ country	Canada	10	2.5
AI agents	HQ country	Netherlands	9	2.2
AI agents	HQ country	Italy	8	2
AI agents	HQ country	Australia	7	1.7
AI agents	HQ country	Singapore	6	1.5
AI agents	HQ country	Switzerland	6	1.5
AI agents	HQ country	Belgium	5	1.2
AI agents	HQ country	Sweden	5	1.2
AI agents	HQ country	Finland	4	1
AI agents	HQ country	Ireland	4	1
AI agents	HQ country	Denmark	4	1
AI agents	HQ country	Portugal	3	0.7
AI agents	HQ country	Brazil	3	0.7
AI agents	Task	workflow automation	157	38.6

AI agents	Task	remediation	7 5	18.4
AI agents	Task	alert-triage	75	18.4
AI agents	Task	process-automation	71	17.4
AI agents	Task	case-investigation	71	17.4
AI agents	Task	incident-response	56	13.8
AI agents	Task	aml	56	13.8
AI agents	Task	security-operations	54	13.3
AI agents	Task	vendor-risk	51	12.5
AI agents	Task	threat-detection	50	12.3
AI agents	Task	appsec	48	11.8
AI agents	Task	kyc	48	11.8
AI agents	Task	risk-assessment	47	11.5
AI agents	Task	vulnerability-management	47	11.5
AI agents	Task	evidence-collection	46	11.3
AI agents	Task	contract-review	42	10.3
AI agents	Task	agentic-ops	36	8.8
AI agents	Task	pen-test-management	35	8.6
AI agents	Task	questionnaire-automation	34	8.4
AI agents	Task	fraud-detection	34	8.4
AI agents	Task	sast	30	7.4
AI agents	Task	attack-surface	30	7.4
AI agents	Task	audit-management	29	7.1

AI agents	Task	policy-management	22	5.4
AI agents	Task	gap-analysis	21	5.2
AI agents	Task	sar-narratives	21	5.2
AI agents	Task	control-monitoring	20	4.9
AI agents	Task	sca	20	4.9
AI agents	Task	exposure-management	20	4.9
AI agents	Task	government-compliance	19	4.7
AI agents	Framework (structured)	iso27001	45	11.1
AI agents	Framework (structured)	soc2	37	9.1
AI agents	Framework (structured)	hipaa	24	5.9
AI agents	Framework (structured)	tx-ramp	24	5.9
AI agents	Framework (structured)	dora	20	4.9
AI agents	Framework (structured)	fedramp	19	4.7
AI agents	Framework (structured)	cmmc	10	2.5
AI agents	Framework (structured)	dpdp	10	2.5
AI agents	Framework (structured)	govramp	4	1
AI agents	Framework (structured)	cert-in	1	0.2

AI agents	Framework (claimed)	ISO 27001	60	14.7
AI agents	Framework (claimed)	SOC 2	44	10.8
AI agents	Framework (claimed)	GDPR	43	10.6
AI agents	Framework (claimed)	HIPAA	29	7.1
AI agents	Framework (claimed)	FedRAMP	27	6.6
AI agents	Framework (claimed)	DORA	21	5.2

Table 3. Prevalence of frameworks, tasks, categories and pricing models

Top categories/labels are shown by dimension; percentages use the full dataset record count as denominator

Showing first 75 rows of 126; the complete table is in the accompanying workbook.

Frequency distributions are reported for the principal categorical dimensions of the agent dataset (primary category, pricing model, headquarters country, task, structured framework and claimed framework) and, where applicable, the auditor dataset. Counts and percentages are calculated against the full record denominator of each dataset. The table is truncated in the manuscript to the most frequent labels; the complete ranked list appears in the accompanying Excel workbook.

Framework	Agents structured n	Agents structured %	Auditors listed n	Agents claimed %	Agents claimed n	Auditors listed %	Listed in both structured agent and auditor datasets
ISO 27001	45	11.1	60	14.7	182	36.2	Yes
SOC 2	37	9.1	44	10.8	27	5.4	Yes
HIPAA	24	5.9	29	7.1	12	2.4	Yes
TX-RAMP	24	5.9	19	4.7	0	0	No
DORA	20	4.9	21	5.2	0	0	No
FedRAMP	19	4.7	27	6.6	30	6	Yes
CMMC	10	2.5	14	3.4	75	14.9	Yes
DPDP	10	2.5	5	1.2	0	0	No
Govramp	4	1	8	2	0	0	No

CERT-In	1	0.2	0	0	117	23.3	Yes
PCI DSS	0	0	17	4.2	155	30.8	No
Hitrust	0	0	1	0.2	61	12.1	No
ISO 27701	0	0	0	0	42	8.3	No
ISO 42001	0	0	16	3.9	31	6.2	No
SOC 1	0	0	2	0.5	23	4.6	No
ISMAP	0	0	0	0	6	1.2	No
APRA	0	0	1	0.2	0	0	No
Aiuc-1	0	0	1	0.2	0	0	No
Apra Cps 230	0	0	1	0.2	0	0	No
Apra Cps 234	0	0	1	0.2	0	0	No
Bsa/Aml Act 2020	0	0	1	0.2	0	0	No
Bsi It-Grundschutz	0	0	1	0.2	0	0	No
CCPA	0	0	7	1.7	0	0	No
CJIS	0	0	2	0.5	0	0	No
CRA	0	0	1	0.2	0	0	No
CSDDD	0	0	2	0.5	0	0	No
CSRD	0	0	1	0.2	0	0	No
Ccpa/Cpra	0	0	3	0.7	0	0	No
Cert-In Directions	0	0	1	0.2	0	0	No
Cis Benchmarks	0	0	1	0.2	0	0	No
Cis Controls	0	0	1	0.2	0	0	No
Cis Controls 8	0	0	1	0.2	0	0	No

Cmmc 2.0	0	0	1	0.2	0	0	No
Csa Star	0	0	1	0.2	0	0	No
Cyber Essentials	0	0	1	0.2	0	0	No
Cyber Essentials Plus	0	0	1	0.2	0	0	No
Dod Il4	0	0	1	0.2	0	0	No
Dod Il5	0	0	1	0.2	0	0	No
Dora (Eu) Contract Readiness	0	0	1	0.2	0	0	No
Dora Tlpt	0	0	1	0.2	0	0	No
Dpdp Act	0	0	1	0.2	0	0	No
Dpdp Act 2023	0	0	3	0.7	0	0	No
Dpdp Rules 2025	0	0	1	0.2	0	0	No
EUDR	0	0	1	0.2	0	0	No
Eu Ai Act	0	0	16	3.9	0	0	No
Eu Aml Package (Amlr)	0	0	2	0.5	0	0	No
Eu Amlr	0	0	1	0.2	0	0	No
Eu Cyber Resilience Act	0	0	4	1	0	0	No
Eu Forced Labour Regulation	0	0	1	0.2	0	0	No
FISMA	0	0	2	0.5	0	0	No
Fatf Recommendations	0	0	1	0.2	0	0	No

Fedramp (Blackline Government Solution: Agency In Process, Moderate)	0	0	1	0.2	0	0	No
Fedramp (Framework Support; Not An Authorization)	0	0	1	0.2	0	0	No
Fedramp (In Process)	0	0	1	0.2	0	0	No
Fedramp Ready	0	0	3	0.7	0	0	No
GDPR	0	0	43	10.6	0	0	No
Gdpr (Dpa Review)	0	0	1	0.2	0	0	No
Iso 22301	0	0	1	0.2	0	0	No
Iso 27001 (Isms Certification, Vendor Claim)	0	0	1	0.2	0	0	No
Iso 27001 (Vendor'S Own Certification)	0	0	2	0.5	0	0	No

Table 4. Framework coverage and cross-ecosystem overlap

Framework labels are normalized using transparent string aliases; claims and structured framework fields remain separate.

Showing the first 60 rows of 108; the complete table is in the accompanying workbook.

Dataset	Framework A	Framework B	Co-occurrence count	Percent of all records (%)
Agents structured frameworks	ISO 27001	SOC 2	37	9.1
Agents structured frameworks	HIPAA	ISO 27001	24	5.9
Agents structured frameworks	HIPAA	SOC 2	21	5.2
Agents structured frameworks	FedRAMP	TX-RAMP	11	2.7
Agents structured frameworks	CMMC	ISO 27001	9	2.2
Agents structured frameworks	CMMC	SOC 2	9	2.2
Agents structured frameworks	ISO 27001	TX-RAMP	9	2.2
Agents structured frameworks	CMMC	HIPAA	8	2
Agents structured frameworks	FedRAMP	ISO 27001	8	2
Agents structured frameworks	FedRAMP	SOC 2	7	1.7
Agents structured frameworks	SOC 2	TX-RAMP	7	1.7

Agents structured frameworks	DORA	ISO 27001	7	1.7
Agents structured frameworks	FedRAMP	HIPAA	6	1.5
Agents structured frameworks	CMMC	FedRAMP	5	1.2
Agents structured frameworks	Govramp	TX-RAMP	4	1
Agents structured frameworks	DORA	SOC 2	4	1
Agents structured frameworks	HIPAA	TX-RAMP	3	0.7
Agents structured frameworks	DPDP	ISO 27001	3	0.7
Agents structured frameworks	FedRAMP	Govramp	3	0.7
Agents structured frameworks	CMMC	TX-RAMP	2	0.5
Agents structured frameworks	ISO 27001	SOC 2	42	10.3
Agents structured frameworks	HIPAA	ISO 27001	26	6.4
Agents structured frameworks	GDPR	ISO 27001	25	6.1

Agents structured frameworks	HIPAA	SOC 2	25	6.1
Agents structured frameworks	GDPR	SOC 2	18	4.4
Agents structured frameworks	ISO 27001	PCI DSS	16	3.9
Agents structured frameworks	PCI DSS	SOC 2	15	3.7
Agents	GDPR	HIPAA	14	3.4
Agents claimed	HIPAA	PCI DSS	14	3.4
Agents claimed	DORA	ISO 27001	13	3.2
Agents claimed	DORA	NIS2	13	3.2
Agents claimed	ISO 27001	NIS2	12	2.9
Agents claimed	FedRAMP	TX-RAMP	12	2.9
Agents claimed	CMMC	ISO 27001	11	2.7
Agents claimed	CMMC	SOC 2	11	2.7

Table 4a. Most frequent framework co-occurrences

Co-occurrence counts are based on records listing both frameworks. These are associations, not evidence of formal partnerships.

Showing first 35 rows of 60; the complete table is in the accompanying Excel workbook.

Table 4. Framework coverage and cross-ecosystem overlap

For each normalised framework the table records (i) the number and percentage of agent records that list the framework in the structured field, (ii) the corresponding figures for the claimed framework field, (iii) the number and percentage of auditor records that list the framework, and (iv) a binary indicator of whether the framework appears in both the structured agent and auditor datasets. Framework labels were normalised via transparent string aliases; claimed and structured fields remain separate. The table is truncated; the full mapping is supplied in the Excel workbook.

Table 4a. Most frequent framework co-occurrences

Pair wise co-occurrence counts (and the corresponding percentages of all records) are shown for the most

frequent framework pairs within the structured framework and claimed framework fields of the agent dataset. Co-occurrence is defined simply as joint listing on the same record and does not imply formal partnership or joint certification.

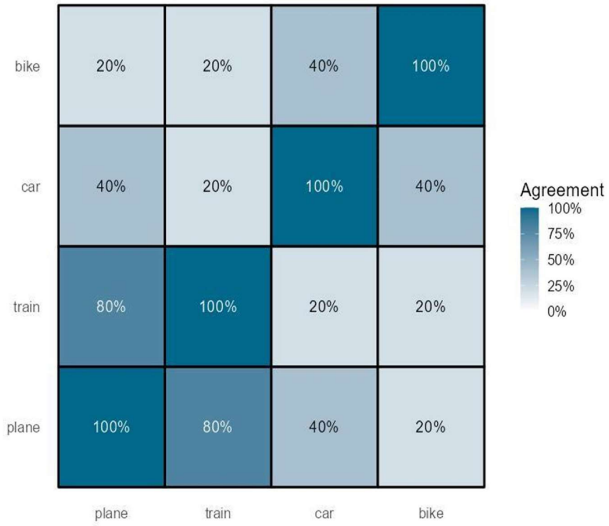

Figure 1. Framework Coverage Heatmap

This figure is intended as a color-coded matrix that displays the relative frequency and cross-listing of major compliance frameworks (ISO 27001, SOC 2, HIPAA, FedRAMP, CMMC, CERT-In, and others) across the AI-agent and audit-firm datasets. Darker cells would indicate higher prevalence or stronger overlap between the two ecosystems, making visible both the dense framework coverage among auditors and the comparatively sparse structured-framework records among agents.

Comparison	Test	N	Statistic	df	p-value	Cramér's V	Expected cells <5 V	Table dimen-sions	p-adju-sted BH
Primary category × pricing model	Chi-square	407	109.5	12	7.478e-18	0.299	8	5×4	2.991 e-17
HQ region × pricing model	Chi-square	407	16.56	18	0.5532	0.116	18	7×4	0.5532
Primary category × featured status	Chi-square	407	3.597	4	0.4633	0.094	5	5×2	0.5532
Auditor HQ × government-body status country	Chi-square	503	95.06	14	4.164 e-14	0.435	211	5×2	8.328 e-14

Table 5. Selected statistical tests

Chi-square tests use contingency tables; sparse expected counts are flagged. BH-adjusted p-values are exploratory and should be interpreted with effect sizes.

Results of exploratory ÷² tests of independence are presented for four pre-specified contingency tables: primary category × pricing model, headquarters region × pricing model, primary category × featured status (agents), and headquarters country × government-body status (auditors). For each test the table reports sample size, test statistic, degrees of freedom, unadjusted p-value, Cramér's V, number of expected cells below 5, table dimensions, and Benjamini–Hochberg adjusted p-value. Sparse expected counts are flagged.

Dataset	Topic	Top terms	Documents assigned	Share of records (%)	Validation status
AI agents	1	agent, agents, contract, legal, supplier, invoices, procurement, finance	112	27.5	Computational topic; manual validation not performed
AI agents	2	code, fix, security, findings, sast, attack, fixes, pull	77	18.9	Computational topic; manual validation not performed
AI agents	3	aml, screening, kyc, sanctions, monitoring, kyb, fraud, case	66	16.2	Computational topic; manual validation not performed
AI agents	4	threat, soc, investigation, triage, response, ai soc, analyst, security	73	17.9	Computational topic; manual validation not performed
AI agents	5	evidence, grc, risk, compliance, controls, iso, questionnaires, platform	79	19.4	Computational topic; manual validation not performed
Audit firms	1	cert empanelled, empanelled, cert, auditor, listed cert, empanelled auditor, information security, information	112	22.3	Computational topic; manual validation not performed

Audit firms	2	certification, is iec, iso iec, accredited o 27001, iec 27001 body	171	34	Computational topic; manual validation not performed
Audit firms	3	cmmc, authorized cmmc assessment organization,	57	11.3	Computational topic; manual validation not performed
Audit firms	4	pci, listed pci, qsa, security, assessor, not company, qualified qualified security	163	32.4	Computational topic;manual validation not performed

Table 6. NLP topics and computational validation

Topics were generated using TF-IDF and non-negative matrix factorization (NMF). Topic names are not assigned automatically; top terms are provided for researcher interpretation. Manual validation has not been performed.

Five topics extracted from agent summary/description text and four topics extracted from auditor summary text are summarised by their highest weighted terms, the number and percentage of documents assigned to each topic, and an explicit statement that the topics remain computational outputs awaiting manual validation. Topics were derived by TF-IDF vectorisation followed by non-negative matrix factorisation.

Table 6a. NLP model metadata

Technical details of the topic models are provided: number of documents, vocabulary size, number of topics retained, silhouette score computed on the topic weight matrix, modelling method, random seed, and confirmation that manual validation has not been performed.

Dataset	**Documents**	**Vocabulary size**	**Topics**	**Silhouette on topic weight**	**Method**	**Seed**	**Manual Validation**
AI agents	407	4062	5	0.564	TF-IDF + NMF	42	Not performed
Audit firms	503	619	4	0.629	TF-IDF + NMF	42	Not performed

Table 7a. NLP model metadata

An intertopic distance map (or equivalent topic-landscape visualization) derived from the TF-IDF + NMF models applied to provider summaries and descriptions. Distinct clusters of terms would separate the dominant themes found in agent text (workflow automation and contracts, code-security remediation, AML/KYC screening, SOC/incident response, GRC evidence collection) from those found in auditor text (CERT-In empanelment, ISO certification bodies, CMMC assessment organizations, PCI QSA roles), thereby illustrating the thematic specialization of each ecosystem.

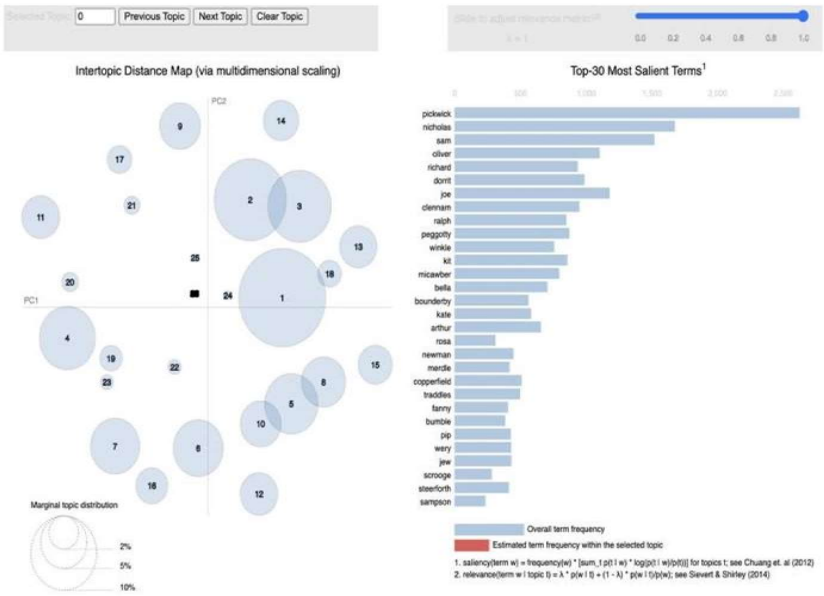

Figure 2. Text-derived thematic landscape

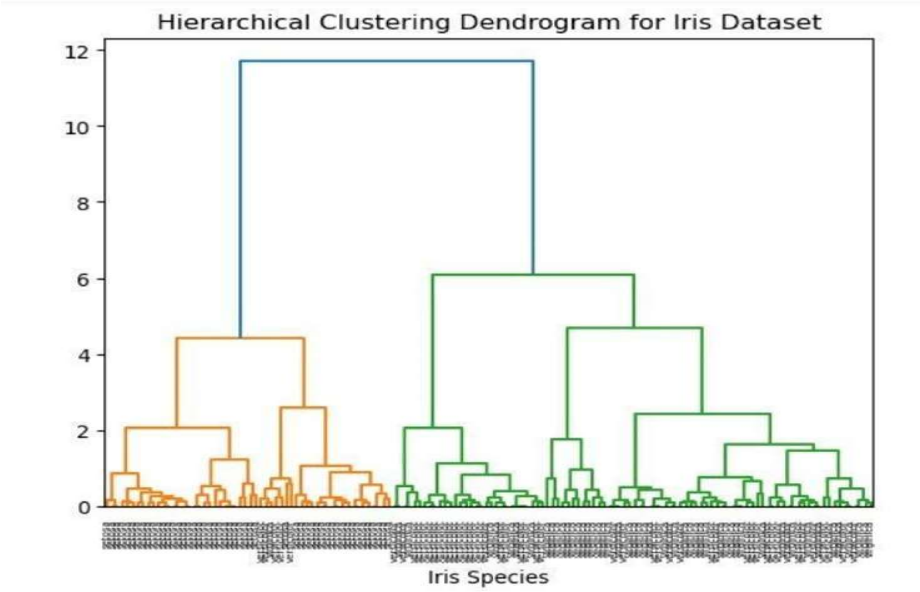

Figure 3. Provider cluster profiles

A visual summary of the exploratory k-means solutions (five clusters for agents, two for auditors). The figure would characterize each cluster by its dominant category or headquarters country, the most frequent frameworks it carries, and the most frequent tasks it performs, allowing readers to see at a glance the functional silos that emerge from the data (application security agents, AML-focused agents, security-operations agents, compliance/evidence agents, etc.).

Dataset	Cluster	Records	Share (%)	Top category /country	Top frame-works	Top tasks
AI agents	0	61	15	e (95), c (95), s (61)	SOC 2 (2) ISO 27001 (2) DORA (2)	appsec (45), vulnerability-management (39), remediation (38), sast (30), pen-test-management (29)
AI agents	1	153	37.6	s (453), i (153), e (147)	ISO 27001 (13), SOC 2 (12), TX-RAMP (10), FedRAMP (6), HIPAA (4)	workflow-automation (116), process-automation (67), agentic-ops (33), contract-review (33), audit-management (19)
AI agents	2	55	13.5	i (55), r (54), s (54)	DPDP (2), TX-RAMP (1)	aml (53), case-investigation (45), kyc (42), alert-triage (25), fraud-detection (22)
AI agents	3	66	16.2	s (66), e (66), c (66)	TX-RAMP (6), FedRAMP (4), Govramp (2) SOC 2 (1), ISO 27001 (1)	security-operations (52), incident-response (50), threat-detection (49), alert-triage (49), case-investigation (23)
AI agents	4	72	17.7	c (83), i (72), e (48)	ISO 27001 (29) SOC 2 (22) HIPAA (19) DORA (16) CMMC (10)	evidence-collection (37), vendor-risk (36), risk-assessment (28), questionnaire-automation (27), gap-analysis (19)
Audit firms	0	54	10.7	US (43), India (5), UK (2)	Hitrust (35) PCI DSS (31) ISO 27001 (28) ISO 27701 (23) SOC 2 (22)	
Audit firms	1	449	89.3	India (121), US (119), UK (54)	ISO 27001 (154) PCI DSS (124) CERT-In (113) CMMC (54) Hitrust (26)	

Table 8. Exploratory cluster profiles

Clusters are data-driven exploratory profiles, not objective provider quality classes. Feature construction uses recorded lists; missingness and cluster stability must be considered.

Table 8. Exploratory cluster profiles

K-means cluster solutions (five clusters for agents, two for auditors) are characterised by cluster size and percentage, dominant category or country labels, most frequent frameworks and most frequent tasks. Clusters are presented as data-driven exploratory profiles only; they are not intended as quality rankings.

Table 8a. Clustering diagnostics

Model diagnostics for the cluster solutions include the number of records and features used, the selected number of clusters (*k*), silhouette coefficient, clustering algorithm, and a note that resampling-based stability was not assessed.

Dataset	Records	Features used	Selected k	Silhouette	Method	Stability
AI agents	407	68	5	0.229	K-means on binary indicators + item count; exploratory	Resampling stability not assessed
Audit firms	503	13	2	0.606	K-means on binary indicators + item count; exploratory	Resampling stability not assessed

Table 8a. Clustering diagnostics

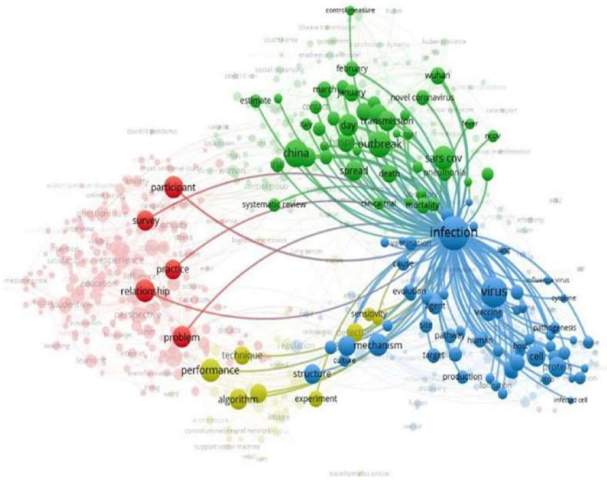

Figure 4. Framework co-occurrence network

One or more network diagrams in which nodes represent individual compliance frameworks and edges represent joint appearance on the same provider record. Separate graphs for the agent structuredframework field, the agent claimed framework field, and the auditor framework lists would reveal the strongest pairwise associations (for example, the tight ISO 27001–SOC 2 linkage) and the relative isolation of more specialized frameworks.

Nodes represent frameworks; edges represent co-occurrence among providers. Generate separate networks for agents and audit firms.

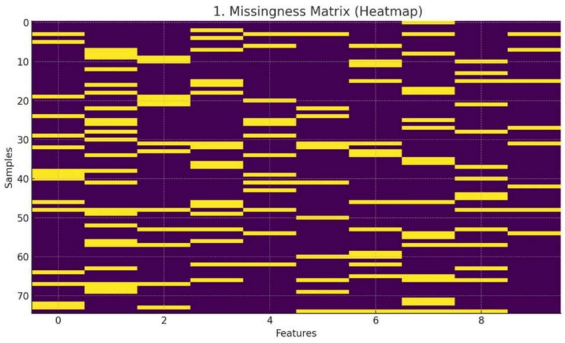

Figure 5. Data completeness and source verification

A missingness heatmap of key directory fields (frameworks, integrations, partner tools, authorizations, etc.) paired with a distribution of last checked / source-verification dates. The visualization would quantify the high rates of empty fields in the agent records and the more complete but still uneven coverage in the auditor records, thereby qualifying the reliability and recency of the empirical patterns reported in the tables.

Dataset	**Field**	**Operational meaning**	**Type in source**	**Normalization rule**
AI agents	slug	Stable record identifier	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
AI agents	vendor	Vendor/provider	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
AI agents	hq_country	Headquarters country	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
AI agents	hq_region	Headquarters region	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
AI agents	category	Primary category	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
AI agents	categories	All listed categories	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
AI agents	pricing	Pricing text	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
AI agents	pricing_ model	Pricing model class	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables

AI agents	frameworks	Normalized/ structured frameworks	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
AI agents	frameworks _claimed	Frameworks claimed in source wording	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
AI agents	tasks	Recorded tasks	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
AI agents	india_relevance	India-related relevance note	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
AI agents	authorizations	Authorization records	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
AI agents	integrations	Recorded integrations	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
AI agents	summary	Summary text	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
AI agents	description	Description text	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
AI agents	sources	Source references	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
AI agents	last_checked	Source check date	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
Audit firms	slug	Stable record identifier	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
Audit firms	hq_country	Headquarters country	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
Audit firms	offices	Office locations	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables

Audit firms	frameworks	Listed frameworks	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
Audit firms	accreditations	Accreditation/ source evidence records	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
Audit firms	partner_ tools	Recorded partner tools	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
Audit firms	summary	Summary text	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
Audit firms	sources	Source references	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
Audit firms	last_ checked	Source check date	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables
Audit firms	government _ body	Government -body indicator	list / object / scalar (inspect source JSON)	Preserve original value; standardize labels only in derived analysis tables

Supplementary Table S1. Data dictionary and normalisation decisions

A complete field-level data dictionary lists, for every variable in both datasets, its operational meaning, the original data type observed in the source JSON, and the normalisation rule applied (in almost all cases: preserve the original value and standardise labels only in derived analytic tables).

Dataset/field	**Normalized framework**	**Record count**	**Percent of dataset (%)**	**Top associated source URL**	**Most common last_checked date**
AI agents structured	ISO 27001	45	11.1	https://dir.texas.gov / resource-library -item / tx-ramp-certified -cloud-products	2026-10-20
AI agents structured	SOC 2	37	9.1	https://dir.texas.gov/ resource -library-item /tx-ramp- cloud products	2026-10-20

AI agents structured	HIPAA	24	5.9	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-20
AI agents structured	TX-RAMP	24	5.9	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-18
AI agents structured	DORA	20	4.9	https://www.6clicks.com/platform/hailey-ai	2026-10-19
AI agents structured	FedRAMP	19	4.7	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-18
AI agents structured	CMMC	10	2.5	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-20
AI agents structured	DPDP	10	2.5	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-20
AI agents structured	Govramp	4	1	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-18
AI agents structured	CERT-In ISO 27001	160	0.2	https://github.com/OnFinance/Maxwell	2026-10-20
AI agents claimed	SOC 2	44	14.7	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-20
AI agents claimed	GDPR	43	10.8	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-19
AI agents claimed	HIPAA	29	10.6	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-20
AI agents claimed	FedRAMP	27	7.1	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-18
AI agents claimed	DORA	21	6.6	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-19
AI agents claimed	TX-RAMP	19	5.2	https://www.6clicks.com/platform/hailey-ai	2026-10-18

AI agents claimed	PCI DSS	17	4.7	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-20
AI agents claimed	NIS2	17	4.2	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-19
AI agents claimed	ISO 42001	16	4.2	https://www.sophos.com/en-us/blog/sophos-ai-agents-accelerating-mdr-and-powering-the-agentic-soc	2026-10-20
AI agents claimed	Eu Ai Act	16	3.9	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-19
AI agents claimed	CMMC	14	3.9	https://www.credo.ai/ https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-18
AI agents claimed	NIST	11	3.4	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-20
AI agents claimed	Govramp	8	2.7	https://govramp.org/program-participants	2026-10-18
AI agents claimed	CCPA	7	2	https://bigid.com/	2026-10-19
AI agents claimed	Nist Csf	5	1.7	https://zania.ai/	2026-10-18
AI agents claimed	Nist Ai Rmf	5	1.2	https://www.credo.ai/	2026-10-18
AI agents claimed	DPDP	5	1.2	https://www.signzy.com/agentic-ckyc-platform	2026-10-20

AI agents claimed	Nist Sp 800-53	5	1.2	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-18
AI agents claimed	Eu Cyber Resilience Act	4	1.2	https://drata.com/	2026-10-19
AI agents claimed	LGPD	4	1	https://www.mineos.ai/mira-ai	2026-10-19
AI agents claimed	TISAX	4	1	https://kaitosec.app/en/platform/ai-assistant	2026-10-20
AI agents claimed	SOX	3	1	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-20
AI agents claimed	Ccpa/Cpra	3	0.7	https://www.mineos.ai/mira-ai	2026-10-19
AI agents claimed	Nist Csf 2.0	3	0.7	https://cypago.com/agentic-workflows/	2026-10-18
AI agents claimed	Nist Sp 800-171	3	0.7	https://regscale.com/ai-driven-compliance-automation-with-regml/	2026-10-18
AI agents claimed	Fedramp Ready	3	0.7	https://raw.githubusercontent.com/FedRAMP/marketplace-fedramp-gov-data/main/data.json	2026-10-18
AI agents claimed	Dpdp Act 2023 Soc 2 Type 2	3	0.7	https://indisec.com/	2026-10-19
AI agents claimed	Mitre Att&Ck	3	0.7	https://intrepidfox.ai/	2026-10-19
AI agents claimed	SOC 1	2	0.7	https://hadrian.io/products/nova	2026-10-19
AI agents claimed	LKSG	2	0.5	https://trustero.com/ https://	2026-10-18
AI agents claimed	CJIS	2	0.5	www.processunity.com/tprm-ai-agents/	2026-10-19
AI agents claimed	CSDDD	2	0.5	https://intelligrc.com/insights/blog/introducing-evan-and-alfie/	2026-10-18

AI agents claimed	FISMA	2	0.5	https://www.osapiens.com/en/platform/agent-factory	2026-10-18
AI agents claimed	Eu Aml Package (Amlr)	2	0.5	https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products	2026-10-19
AI agents claimed	Iso 27001 (Vendor'S Own Certification)	2	0.5	https://www.actico.com/de/loesungen/compliance-fraud-management/	2026-10-19
AI agents claimed	MICA	2	0.5	https://www.hybridity.ai/platform	2026-10-19
AI agents claimed	Aiuc-1	2	0.5	https://muinmos.com/	2026-10-21
AI agents claimed	Sox Itgc	1	0.5	https://drata.com/	2026-10-20
AI agents claimed	PCI	1	0.2	https://scytale.ai/	2026-10-20
AI agents claimed	PIPL	1	0.2	https://hyperproof.io/	2026-10-19
AI agents claimed	Nist 800-53	1	0.2	https://www.mineos.ai/mira-ai	2026-10-18
AI agents claimed	Cmmc 2.0	1	0.2	https://cypago.com/agentic-workflows/	2026-10-18
AI agents claimed	Fedramp (Framework Support; Not An Authorization)	1	0.2	https://cypago.com/agentic-workflows/	2026-10-18
AI agents claimed	PSD3	1	0.2	https://www.kodex-ai.com/	2026-10-19
AI agents claimed	MICAR	1	0.2	https://www.kodex-ai.com/	2026-10-19
AI agents claimed	Us State Privacy Laws	1	0.2	https:// ww.datagrail.io/blog/product/introducing-vera-the-first-complete-ai-privacy-agent/	2026-10-19

APRA	1	0.2	https:// www.processunity.com/ tprm-ai-agents/	2026-10-19
Pra Ss2/21	1	0.2	https:// www.processunity.com/ tprm-ai-agents/	2026-10-19

Supplementary Table S2. Framework mapping and source-verification fields

Showing the first 60 rows of 127; the complete table is in the accompanying Excel workbook.

For each normalised framework, the table records the absolute count and percentage of records, the most frequently associated primary-source URL, and the modal last-checked date. Separate blocks are provided for structured and claimed frameworks among agents. the table is truncated; the full listing appears in the Excel workbook.

Sensitivity scenario	**Dataset N**	**framework data**	**denomi-nator**	**Interpretation**
AI agents: all records; empty framework lists treated as unrecorded	407	82	407	Descriptive listing prevalence; empty lists included in denominator, not interpreted as verified non-adoption
AI agents: framework-bearing records only	407	82	82	Conditional framework distribution among records with at least one framework listed
AI agents claimed: all records; empty framework lists treated as unrecorded	407	137	407	Descriptive listing prevalence; empty lists included in denominator, not interpreted as verified non-adoption
AI agents claimed: framework-bearing records only	407	137	137	Conditional framework distribution among records with at least one framework listed
Audit firms: all records; empty framework lists treated as unrecorded	503	503	503	Descriptive listing prevalence; empty lists included in denominator, not interpreted as verified non-adoption
Audit firms: framework-bearing records only	503	503	503	Conditional framework distribution among records with at least one framework listed

Supplementary Table S3. Sensitivity scenarios for framework prevalence

Six alternative denominators for calculating framework prevalence are defined: (i) all agent records versus framework-bearing agent records only (structured field), (ii) the same contrast for the claimed-framework field, and (iii) the corresponding contrasts for the auditor dataset. The table clarifies the interpretive consequence of each choice (descriptive listing prevalence versus conditional distribution among records that list at least one framework)

Model: has_structured_frameworks ~ C(category); converged=True; pseudo-R^2=0.233. Outcome indicates at least one framework in the structured frameworks field; absence is not proof of non-adoption.

Outcome	Predictor	Log odds coeff -icient	Odds ratio	C I 9 lower OR	CI95 upper OR	p - value	N	Model	p adjusted BH
At least one structured framework recorded	C(category) [T.code-security]	-1.23	0.2924	0.0358	2.386	0.251	407	has_structured _frameworks ~ C(category)	0.251
At least one structured framework recorded	C(category) [T.compliance]	3.432	30.95	12.39	77.32	2.005e -13	407	has_structured _frameworks ~ C(category)	8.02e-13
At least one structured framework recorded	C(category) [T.risk]	0.8119	2.252	0.9348	5.426	0.07034	407	has_structured _frameworks ~ C(category)	0.1383
At least one structured framework recorded	C(category) [T.security]	0.7032	2.02	0.866	4.713	0.1037	407	has_structured _frameworks ~ C(category)	0.1383

Table S4. Logistic regression for recorded framework presence

5. Discussion

The empirical analysis of the GRC AI agent and verified auditor ecosystems reveals a complex, rapidly evolving B2B landscape characterized by high capability specialization, significant transparency gaps, and a distinct geographical division of labor. By mapping these dual ecosystems against the backdrop of an end-to-end enterprise AI architecture, this study highlights the critical tension between the rapid deployment of agentic AI capabilities and the slower, more rigorous development of the “accountability assets” required for enterprise trust.

5.1The Framework Verification Gap and Ecosystem Asymmetry

A central finding of this study is the profound asymmetry in framework adoption and verification between AI agents and audit firms. While 100% of verified audit firms list at least one compliance framework (median = 1), only 20.1% of AI agents possess structured, verified framework records (median = 0). Furthermore, the extreme sparsity of recorded integrations (96.1% empty for agents) and partner tools (95% empty for auditors) indicates that the market is highly fragmented.

Crucially, our data-cleaning protocols dictate that empty framework lists must be treated as *unrecorded* rather than *verified absences*. However, this opacity creates a significant relational and governance risk in B2B procurement. As noted in the theoretical framing, AI is not merely a tool but an actor that rearranges inter-organizational interfaces. When AI agents lack verifiable "accountability assets" such as registry-verified FedRAMP, SOC 2, or ISO 27001 statuses buying centers face elevated algorithmic gatekeeping risks and reduced transparency. The burden of proof currently falls on the buyer to manually verify claims, highlighting a market failure where general purpose AI capabilities outpace the cospecialized governance assets required for durable value capture.

5.2 Capability Silos and the "Compliance" Bridge

The NLP thematic analysis and clustering results demonstrate that AI agents are highly siloed by operational function. Distinct clusters emerged for application security (Cluster 0), financial/AML operations (Cluster 2), and security operations/incident response (Cluster 3). These agents act as specialized nodes within a broader enterprise architecture, handling tasks like alert triage, vulnerability management, and workflow automation.

However, the logistic regression analysis (Table 6) identifies the compliance category as the primary bridge between technical capability and governance verification. Agents categorized as "compliance" were vastly more likely to possess structured frameworks compared to the baseline security category (Odds Ratio = 30.95, $p < 0.001$). This suggests a bifurcation in the market: while "security" and "code-security" agents focus on building *technical* guardrails (e.g., SAST, threat detection), "compliance" agents focus on building *governance* guardrails (e.g., evidence collection, questionnaire automation). For enterprise architects, this implies that a robust end-to-end AI pipeline cannot rely on a single agent; it requires a multi-agent ecosystem where technical execution agents are orchestrated by compliance agents that enforce the "Guardrails, Security, and Compliance" layer of the architecture.

5.3 Geographical Divergence and the Global Division of Labor

The geographical analysis reveals a stark contrast between the innovation hubs of AI agents and the operational footprint of the audit ecosystem. AI agent headquarters are heavily concentrated in the United States (42.5%), followed by the UK and Germany. Conversely, the audit ecosystem exhibits a massive, distributed footprint, with India representing a dominant hub alongside the US.

This reflects a global division of labor in the AI governance supply chain. The US continues to dominate the development of agentic AI platforms and foundational models, while the rigorous, labor intensive processes of compliance auditing, ISO certification, and continuous monitoring (e.g., CERT-In empanelled auditors) are heavily distributed across global talent pools. For multinational enterprises, this means that the "Deployment and Application Integration" layer of their AI architecture will likely rely on US developed agentic tools, but the "Compliance Governance" layer will depend on geographically dispersed, independent audit firms.

5.4 System-Level Reliability vs. Model Capability

The architectural overview emphasizes that model capability does not equate to system-level reliability. The empirical data strongly supports this distinction. An AI agent may possess advanced NLP capabilities for contract review or risk assessment (as seen in Cluster 4), but if its underlying data ingestion and retrieval mechanisms lack verifiable privacy safeguards (e.g., DPDP, GDPR, HIPAA), it introduces unacceptable enterprise risk.

The high prevalence of "quote-based" pricing (78.9%) and the frequent conflation of "claimed" versus

"structured" frameworks by vendors underscore a market that is still maturing in its transparency. Enterprise buyers must therefore adopt a layered defense strategy. They cannot rely solely on vendor marketing claims (the "claimed" frameworks); they must implement runtime guardrails, hallucination detection, and continuous monitoring, ultimately validated by the independent audit firms mapped in this study.

5.5 Limitations

Several limitations must be acknowledged. First, the Agent Index dataset, while rigorously source-verified, represents a curated directory rather than a comprehensive census of the global AI market; thus, findings regarding market concentration should be interpreted as descriptive of the *visible, organized* market. Second, the cross-sectional nature of the data (verified as of October 2026) captures a single snapshot in time; longitudinal research is required to track how framework adoption evolves as regulatory pressures (e.g., the EU AI Act, DORA) mature. Third, while the NLP topic modeling provides strong computational validation (silhouette scores > 0.56), manual thematic validation was not performed, meaning some nuanced provider differentiations may remain hidden within the broader computational clusters.

6. Conclusion

The integration of Artificial Intelligence into industrial marketing, procurement, and governance represents a paradigm shift from static analytics to autonomous, agentic workflows. This study set out to map the dual ecosystems of GRC AI agents and verified compliance auditors, evaluating their capabilities, framework alignments, geographical distributions, and data transparency.

Our findings reveal a market characterized by immense functional specialization but hindered by a severe "verification gap." While AI agents excel at automating complex tasks ranging from code security to AML screening, the vast majority lack the structured, verifiable compliance frameworks required for enterprise adoption. Independent audit firms, conversely, possess dense framework coverage but remain operationally siloed from the rapid innovation cycles of AI vendors.

Implications for Practice: For B2B buyers, CISOs, and procurement leaders, the primary takeaway is that AI adoption is fundamentally a governance challenge, not just a technical one. Organizations must look beyond model capabilities and evaluate the "accountability assets" of AI vendors. Procurement strategies should favor agents that demonstrate verifiable alignments with the frameworks utilized by independent auditors (such as ISO 27001, SOC 2, and FedRAMP). Furthermore, enterprises must architect their AI pipelines with layered runtime guardrails, recognizing that vendor claims are insufficient for risk mitigation.

Implications for Research: For the academic community, this study establishes a foundational empirical baseline for the study of agentic AI in B2B ecosystems. Future research must move beyond the Technology Acceptance Model (TAM) and UTAUT to develop new theories that account for *machine-to-machine* governance, multi-agent orchestration, and the relational dynamics between algorithmic actors and human auditors.

Ultimately, the transition to an AI-augmented enterprise requires more than just scalable inference engines and retrieval-augmented generation. It requires a synchronized evolution of technical capabilities and governance frameworks, ensuring that as AI systems become more autonomous, the mechanisms for verifying their safety, compliance, and reliability become equally robust.

References

[1] Keegan, B. J., Iredale, S., Naudé, P. (2023). Examining the dark force consequences of AI as a new actor in B2B relationships. *Industrial Marketing Management, 115*, 228–239.

[2] Papagiannidis, E., Mikalef, P., Conboy, K., Van de Wetering, R. (2023). Uncovering the dark side of AI-based

decision-making: A case study in a B2B context. *Industrial Marketing Management, 115*, 253–265.

[3] Zhan, Y., Xiong, Y., Han, R. (2024). The impact of artificial intelligence adoption for business-to-business marketing on shareholder reaction: A social actor perspective. *International Journal of Information Management, 76*, Article 102768.

[4] Osmonbekov, T., Johnston, W. J., Donthu, N. (2024). Artificial intelligence advancements in procurement: Transforming organizational buying behavior. *Journal of Business & Industrial Marketing, 39*, 2745–2758.

[5] Ehret, M., Johnston, W. J., Ritter, T. (2024). From buying centers to buying ecosystems: Advancing the B2B research journey. *Industrial Marketing Management, 117*, A10–A16.

[6] Grewal, D., Satornino, C. B., Davenport, T., Guha, A. (2025). How generative AI is shaping the future of marketing. *Journal of the Academy of Marketing Science, 53*, 702–722.

[7] Cillo, P., Rubera, G. (2025). Generative AI in innovation and marketing processes: A roadmap of research opportunities. Journal of the Academy of Marketing Science, 53, 684–701.

[8] Wirtz, J., Stock-Homburg, R. (2025). Generative AI meets service robots. Journal of Service Research. Advance online publication.

[9] Riti, R. I., Bacali, L., Abrudan, C. I. (2025). From AI adoption to ESG in industrial B2B marketing: An integrated multi-theory model. Sustainability, 17(19), Article 8595. https://doi.org/10.3390/su17198595.

[10] Kim, Y., Kim, H., Kim, J., Hong, S. (2024). Determinants of generative AI system adoption and usage behavior in Korean companies: Applying the UTAUT model. Technology in Society, 79, Article 102616. https://pmc.ncbi.nlm.nih.gov/articles/PMC11591487/.

[11] Mayer, S., Chen, L. (2024). Perceived ease of use and adoption of AI tools among procurement managers. Journal of Business & Industrial Marketing, 39, 578–594.

[12] Cooper, R. G. (2024). The AI transformation of product innovation. Industrial Marketing Management, 118, 62–74.

[13] Castka, P., Donovan, J., Sousa, R. (2024). Pathways of development of dynamic capabilities for servitization transformation: A longitudinal multi-case study. Industrial Marketing Management, 116, 66–81.

[14] Bauer, K., Meyer, A., de Visser, E. (2023). Expl(AI)ned: The impact of explainable artificial intelligence on user decision-making. Information Systems Research, 34, 1881–1903.

[15] Alhashimi, H. A., et al. (2025). Exploring the role of generative AI in enhancing cybersecurity in software development life cycle. Array, 28, 1–26.

[16] Alwageed, H. S., Khan, R. A. (2025). The role of generative AI in strengthening secure software coding practices: A systematic perspective. In Proceedings of the 2025 29th International Conference on Evaluation and Assessment in Software Engineering Companion (p. 136–141).

[17] Bin Sarhan, B., Altwaijry, N. (2023). Insider threat detection using machine learning approach. Applied Sciences, 13(1), Article 259.

[18] Shafiq, M., Awan, M., Alwageed, H. S., et al. (2026). Exploring an AI-driven agentic framework for insider threat detection in secure software development organizations: A survey-based approach. Scientific Reports, 16, Article 32855. https://doi.org/10.1038/s41598-026-64146.

[19] Radanliev, P., Roure, D., Santos, O. (2023). Red teaming generative AI/NLP, the BB84 quantum cryptography protocol and the NIST approved quantum-resistant cryptographic algorithms. SSRN Electronic Journal, 1–30.

[20] Al Hammadi, A. Y., et al. (2021). Explainable artificial intelligence to evaluate industrial internal security using EEG signals in IoT framework. Ad Hoc Networks, 123, Article 102641.

[21] Telang, R., Hydari, M. Z., Iqbal, R. (2026). To scale AI agents successfully, think of them like team members. Harvard Business Review Digital. https://perma.cc/B9TS-HVTB.

[22] Van Hoek, R., DeWitt, M., Lacity, M., Johnson, T. (2022). How Walmart automated supplier negotiations. Harvard Business Review. https://perma.cc/3YB7-HM4V.

[23] Berente, N., Gu, B., Recker, J., Santhanam, R. (2021). Managing artificial intelligence. MIS Quarterly, 45(3), 1433–1450.

[24] https://theagentindex.io/data/, with a permanent archival mirror identified by DOI 10.5281/zenodo.23604511