The American Journal of Management and Economics Innovations
ISSN 2693-0811 | Open Access

TYPE Original Research
PAGE NO. 130-140
DOI 10.37547/tajmei/Volume07Issue06-14

Designing a Scalable Knowledge Library for Asset Managers: Hierarchical Taxonomy (Stacks–Categories–Subcategories), Metadata Schemas, and Multi-Fund Content Inheritance

OPEN ACCESS

SUBMITTED 29 April 2025
ACCEPTED 17 May 2025
PUBLISHED 30 June 2025
VOLUME Vol.07 Issue 06 2025

CITATION
Dmitrieva, E. (2026). Designing a Scalable Knowledge Library for Asset Managers: Hierarchical Taxonomy (Stacks–Categories–Subcategories), Metadata Schemas, and Multi-Fund Content Inheritance. The American Journal of Management and Economics Innovations, 7(06), 130–140. https://doi.org/10.37547/tajmei/Volume07Issue06-14

Ekaterina Dmitrieva
MVision Private Equity Advisers New York, USA

Abstract: The paper proposes an integrated approach to designing intelligent knowledge management systems for the asset management domain. Against the backdrop of the accelerated expansion of global assets under management, which reached USD 147 trillion by mid-2025, alongside a parallel deterioration in operating profitability by 19% relative to the 2018 level, the ordering and reproducible structuring of corporate information becomes a critically important condition for business resilience. The analysis centers on the formation of a three-layer hierarchical taxonomy as the foundational framework of a domain model, as well as a methodological shift from document-centric descriptions toward system-oriented metadata schemes that ensure end-to-end interoperability of entities, attributes, and relationships across the data landscape. Substantial attention is devoted to the implementation of object-oriented content inheritance mechanisms, enabling the standardization of reusable knowledge fragments, reducing variance in phrasing, and sustaining interpretive consistency as scale increases.

The impact of changes in SEC regulatory requirements and the codification embedded in ILPA 2.0 standards is examined separately, with emphasis on how these developments reshape data architecture principles and elevate requirements for source traceability, version

control, and demonstrable compliance alignment. Quantitative performance indicators are provided for the deployment of automated mechanisms supporting responses to due diligence questionnaires (DDQ), reflecting the effects of knowledge formalization and metadata anchoring. The results indicate that integrating semantic knowledge libraries can reduce labor inputs and documentation processing time by up to 75%, while simultaneously improving accuracy, response stability, and conformance with compliance-control requirements.

Keywords: asset management, hierarchical taxonomy, metadata schemes, content inheritance, SEC Marketing Rule, ILPA 2.0, DDQ automation, knowledge library, private markets.

Introduction

The global asset management industry in 2024–2025 entered a phase of pronounced structural reconfiguration, which specialized analytical centers describe as the "great convergence." While maintaining a record level of assets under management (AUM), reaching USD 147 trillion as of June 2025 [1], the sector is simultaneously experiencing mounting pressure driven by operating margin compression. Rising client-servicing costs, combined with a more complex regulatory environment, has produced a widening gap between the market-value dynamics of assets and the growth of operating leverage: scale has ceased to translate automatically into a commensurate increase in efficiency [1]. As a result, management companies are forced to rethink both internal information-handling loops and external disclosure practices, because prior models of content governance and evidentiary support no longer appear sufficient.

A key driver of change is the shift of the center of gravity toward private markets and the alternative investments segment. Indicative forecasts suggest that by 2030, revenues from asset management in private markets may reach USD 432.2 billion, contributing more than half of the industry's total revenue [2]. It is also emphasized that revenue per billion dollars of AUM in private markets exceeds that of traditional discretionary management by multiples—estimated at roughly four times [2]. That said, heightened returns come with exceptionally strict investor expectations regarding transparency, comparability, and disclosure depth. Institutional investors increasingly insist on data delivered in formats aligned with ILPA 2.0 standards, where the number of mandatory categories of partnership expenses expands from 9 to 22 [3]. This raises requirements not only for reporting completeness, but also for rigorous classification, unambiguous interpretation, and reproducible calculations.

Market shifts unfold in parallel with a fundamental recalibration of the regulatory landscape, particularly visible in the United States following the adoption of the SEC's updated Marketing Rule (Rule 206(4)-1). Regulatory logic has noticeably moved toward substantiation, meaning the mandatory presence of a verifiable evidentiary basis for each marketing claim [4]. In fiscal years 2024–2025, the SEC demonstrated a stringent enforcement posture, delivering a record USD 8.2 billion in penalties and disgorgement, with a meaningful share of recoveries associated with violations of recordkeeping requirements and the disclosure of performance indicators [5]. Separately, a substantial number of matters—exceeding 70 cases—has been noted in connection with the use of "off-channel" communications, making particularly apparent the need for centralized control over knowledge and messaging flows within an asset management organization [6, 7].

The purpose of the study is to develop and substantiate an architecture for a scalable knowledge library for asset managers, combining a three-layer domain taxonomy (stacks–categories–subcategories), activity-based metadata schemes, and multi-fund content inheritance to increase disclosure reproducibility and compliance robustness in the automation of DDQ/RFP.

The author's hypothesis rests on the assumption that, by transitioning from a document-centric storage model to activity-based metadata anchored to business activities, and by applying governed inheritance of "canonical" statements across funds, it becomes possible to reduce DDQ preparation effort by ~75% while also lowering compliance risk through traceability, versioning, and uniformity of marketing and reporting formulations.

Scientific novelty is defined by the proposition that, for the first time in the asset management domain,

a coherent linkage is advanced—"hierarchical taxonomy + activity-based metadata + object-oriented multi-fund inheritance"—which interprets the knowledge library as an evidence infrastructure (evidence layer) under the requirements of the SEC Marketing Rule and ILPA 2.0, and which sets formalized mechanisms for canonization/override/rollout to scale content without sacrificing auditability.

Materials and Methods

The methodological design of the study is grounded in a systems analysis of industry standards, regulatory and legal documents, and practices of implementing knowledge management systems (KMS) in financial organizations. The empirical basis relied on analytical reports by McKinsey, PwC, and Oliver Wyman, aggregating AUM trends and operating-efficiency indicators across more than 300 asset managers represented in 19 countries [1, 8, 9]. This body of evidence enabled a comparison between macro-level market dynamics and applied effects stemming from process digitalization and the formalization of corporate knowledge.

A substantial block of the research comprised a legal and regulatory analysis of SEC initiatives, including enforcement actions related to investment advisers and new reporting requirements for private funds, which presume mandatory quarterly reporting on performance, fees, and expenses [10]. In parallel, an in-depth review was conducted of the standards of the Institutional Limited Partners Association (ILPA), with particular focus on the Quarterly Reporting Standards Initiative (QRSI) published in January 2025 [11]. These documents were treated not merely as external constraints, but as sources of formalizable requirements for data structure, terminological alignment, disclosure completeness, and the reproducibility of the metrics presented.

The technological dimension of the work is based on a case-study assessment of the SANKOFA program implementation at the African Development Bank, considered a representative example of moving from a document-centric paradigm to an activity-based metadata model [12, 13]. In the area of taxonomy design and inheritance mechanism implementation, architectural patterns of industrial-grade enterprise content management systems (CMS) were analyzed, including Adobe Experience Manager (AEM) and CrafterCMS, as well as specialized AI solutions oriented toward automating DDQ response preparation (Arphie, Loopio, Responsive) [14]. This comparison made it possible to identify universal design principles—content atomicity, stable classification rules, and governed inheritance of attributes—and to align them with compliance-control constraints and auditability requirements.

Results and Discussion

The movement of assets under management throughout 2024 was marked by an unprecedented expansion: the aggregate increase amounted to approximately USD 15 trillion, representing the largest annual growth over the past decade [1]. At the same time, the observed dynamics display a pronounced asymmetry. While passive strategies continue to dominate the equity segment, active management retains comparative advantages and concentrates demand primarily in fixed income and private markets, where the configuration of risk, liquidity, and information asymmetry makes selection and monitoring capabilities decisive [1]. Table 1 presents key statistical parameters and forward-looking estimates that shape the functional and quality requirements for modern knowledge management systems.

Table 1. Key metrics of the asset management industry and forecasts through 2030 (compiled by the author based on [1, 2]).

Indicator	2024/2025 Value	2030 Forecast	CAGR
Global AUM	$147 trillion	$200 trillion	6.2%
Private markets revenues	n/a	$432.2 billion	8.2%
Tokenized funds AUM	$90 billion	$715 billion	41%
Global investable wealth	$345 trillion	$482 trillion	5.7%
Profit decline per $1 billion AUM	-19% (since 2018)	-9% (additional forecast)	n/a

The presented indicators capture an imbalance between the expanding operational scale of asset managers and their ability to convert growth into durable profitability. This "profitability paradox" is associated with the fact that traditional information-handling practices have approached the limits of scalability: as product lines become more complex and the number of interaction channels increases, the cost of preparing, verifying, and maintaining the timeliness of disclosed information rises accordingly [9]. An additional factor is client-base fragmentation: serving increasingly differentiated investor groups—from mass affluent segments to family offices—requires variable reporting formats and a different depth of disclosure, increasing the burden on material-preparation processes and the need to control their internal consistency [2].

Heightened SEC oversight in 2024–2025 established a regime in which virtually any marketing communication is evaluated through an anti-fraud lens and substantiation requirements [4]. Enforcement practice in this period was dominated by three classes of violations. First, unsubstantiated statements were recorded, where marketing materials included formulations not supported by promptly retrievable documentary and computational evidence. Second, the use of hypothetical performance became a target of scrutiny, including the presentation of modeled results without adequate verification of relevance to the target audience and without sufficient disclosure of assumptions and risks. Third, "off-channel" communications were identified: employees' use of personal messengers violated requirements for business-record retention and completeness of records, directly affecting compliance-control boundaries [7].

The combination of these factors drove a shift from viewing the knowledge library as a passive repository to interpreting it as an evidence infrastructure. The required system must ensure not only content retention, but also the fixation of a claim's status, data provenance, versioning, and the links between specific statements and the corresponding investment products, strategies, and reporting metrics. In this context, the ILPA 2.0 standard published in January 2025 sets a higher level of disclosure granularity and comparability: transparency is defined not exclusively at the fund level, but extends to the level of portfolio companies, with the differentiation of expenses across internal and external components [3]. Table 2 summarizes a comparative description of expense categories under ILPA standards.

Table 2. Comparison of expense categories under ILPA standards (compiled by the author based on [3]).

Characteristic	ILPA 2016	ILPA 2.0 (2025)
Number of expense categories	9	22
Personnel expense granularity	Aggregated	Identification of internal chargebacks (Chargebacks)
Credit line reporting	Optional	Mandatory return calculation with/without credit line impact
Fund formation expenses	General	Separation of syndication and placement agent costs
Feeder funds	Minimal	Full transparency at feeder level

The transition to ILPA 2.0 requirements objectively forces asset managers to synchronize accounting loops and knowledge management systems (KMS) in a manner that ensures continuous aggregation of source data for quarterly reporting. In the absence of hierarchically organized metadata, report preparation inevitably shifts into the domain of manual compilation, reconciliation, and repeated validation of metrics, making the process economically unjustifiable as the number of funds, portfolio companies, and expense types grows.

The structural deficit of most corporate knowledge libraries is not tied to the volume of stored materials, but to the absence of a classification model aligned with business-process logic. Folder-based organization—"Marketing," "Legal," "Investments"—reproduces departmental boundaries and creates information silos in which the same semantic object is duplicated in different versions and with different terminology. A scalable solution presupposes the use of a three-level taxonomy that provides stable context for each atomic knowledge fragment and ensures uniform navigation and search [20]. The top level is formed by "stacks" that aggregate knowledge by functional domains: the corporate stack covers information on the firm, governance, and ESG policies; the investment stack fixes strategies, decision-making procedures, and valuation methodologies; the product stack concentrates on the characteristics of specific funds, their legal structures, fees, and performance; the operational stack describes IT infrastructure, business continuity, and compliance controls. The middle level is represented by "categories" that define thematic areas within each stack; for the investment stack, such areas include, for example, idea generation, risk management, and the exit process. The bottom level is formed by "subcategories," mapped to specific investor request types or DDQ sections; these may include illiquid asset valuation policies or the methodology for tracking error calculation [19, 23].

The practical value of such a hierarchy manifests in the ability to rely on machine learning methods for automated classification of incoming materials and subsequent routing of content into relevant contexts. Within a "sandwich approach," top-down classification is applied to identify the stack at the first step, while bottom-up classification—based on lexical-semantic similarity estimation—is used to refine the subcategory [22]. It is shown that taxonomies with a depth of approximately 5–6 levels create conditions for automating the classification of suppliers and services with accuracy comparable to expert judgment, which becomes critically important for spend analytics and expense reporting preparation [21].

Scalability is determined not only by the presence of a hierarchy, but also by the quality of the metadata scheme serving as the "language" for describing knowledge inside the system. Implementation practice

under the SANKOFA program demonstrates that shifting from describing document types to describing business activities materially simplifies content governance and reduces the cost of keeping content current [13]. In the classical model, metadata is assigned to the document as a container—for example, "fund presentation" or "legal opinion." In an activity-based approach, attributes are attached to a process or event—for example, "launch of Fund X" or "quarterly review of Strategy Y"—and any materials placed within the frame of the corresponding activity automatically inherit its properties and contextual parameters [13]. This logic reduces the likelihood of labeling divergence, simplifies end-to-end aggregation for reporting, and forms a basis for reproducible traceability required in compliance reviews and regulatory requests.

Table 3 specifies an activity-level metadata scheme.

Table 3. Activity-level metadata scheme (compiled by the author using the asset management example based on [13]).

Metadata field	Data type	Description / Source
Activity Identifier	Unique ID	Link to master data in SAP/CRM (e.g., Fund_ID)
Strategy Cluster	Controlled vocab	Mapping to the overarching investment stack
Asset Class	Taxonomy link	Classification by asset type (PE, RE, VC, Private Credit)
Regulatory Regime	Multi-select	Applicable legal regimes (UCITS, AIFMD, SEC)
Sensitivity Level	Label	Confidentiality label (Public, Internal, Confidential)
Status	State	Current lifecycle phase (Active, Closed, In-liquidation)

Applying an activity-based approach predictably reduces the number of maintained metadata schemes by shifting descriptive load from the level of individual documents to the level of stable business activities. In the SANKOFA case, a transition was recorded from 1,150 document-level schemes to 150 activity-anchored schemes, materially reducing administrative complexity and the probability of labeling divergence [13]. For an asset manager, such an outcome means that a fund's marketing brochure receives key attributes—"Investment Manager," "Region," "Sector," "Risk Level"—automatically, based on its belonging to the context of the corresponding fund, without the need for manual tag assignment on each material instance. As a result, classification reproducibility is ensured, and the quality of downstream aggregation for reporting and disclosure increases.

In contemporary asset management, a single investment product is often represented by multiple legal and distribution configurations: a master fund in Luxembourg, feeder structures in the United States (Delaware), and separate managed accounts (SMAs) for large clients. Under high wrapper variability, the substantive layer remains largely common: an estimated 80–90% of descriptions coincide, including strategy exposition, team information, and risk-management principles. In this context, a multi-fund content inheritance mechanism becomes critical, eliminating duplication and enabling governance of a single "canonical" set of statements. In systems of the Adobe Experience Manager (AEM) class, this logic is

implemented via the Multi-Site Manager (MSM), which provides a governed hierarchy between the source and derivative objects [15]. A master object (Blueprint) stores the normative version of content; therefore, an edit—for example, the biography of the lead portfolio manager—is executed once in the canonical instance. Derivative representations are formed as "live copies," automatically created for specific funds or distribution channels and maintained in synchronized alignment with the master object. Where necessary, local overrides are permitted, allowing inheritance to be broken for a strictly defined fragment: a U.S.-specific legal disclaimer may be set locally, while the investment-process description remains inherited [17]. Cascading updates propagate master changes to child objects using configurable triggers—from immediate application to delayed publication after local compliance approval [15, 16].

The functional significance of this model is directly tied to the SEC Marketing Rule: when corrections are required in a strategy risk description, the architecture provides assurance that a single change will be consistently reflected across all derivative marketing materials for multiple feeder structures, minimizing the probability of contradictory disclosure and reducing compliance risk [4].

Within such an architecture, the knowledge library serves as the base layer for automating responses to investor questionnaires (DDQ) and requests for proposals (RFP). Beginning in 2024, solutions combining generative AI with verified knowledge repositories under a RAG (Retrieval-Augmented Generation) pattern have become widely adopted: response generation relies on retrieving relevant approved fragments and is accompanied by citation linkage to internal sources, increasing verifiability and reducing the risk of "plausible" yet factually incorrect formulations [14]. Practical effects are expressed through measurable indicators: savings of more than 30 hours of manual work per complex submission are reported, collectively reducing proposal preparation costs by up to 75% [18]. Automation also enables handling roughly 50% more incoming requests without proportionate headcount expansion in investor relations (IR) teams [18]. Increased accuracy due to semantic search is also material: current answers can be retrieved even when an investor's question wording diverges significantly from typical templates, because matching is performed by semantic proximity rather than literal term coincidence [24].

Table 4 provides return-on-investment (ROI) indicators for deploying automated KMS in 2025.

Table 4. Return-on-investment (ROI) indicators for deploying automated KMS in 2025 (compiled by the author based on [18]).

Metric	Traditional approach	KMS- and AI-based approach	Effect
Time to prepare DDQ	12–15 days	2–3 days	-80%
Cost to process one document	$12–15	$3–4	-70%
Win rate	Baseline	Increase by 50%	+50%
Direct ROI per $1 invested	n/a	Up to $746 (Microsoft case)	Record-high
System payback period	n/a	< 12 months (for 74% of companies)	High

A major effect of deployment is compliance-risk reduction through automated control over the

completeness and currency of the evidence base. Platforms in this class can track document and certificate validity windows and match the requirements of a specific request against the available corpus of approved materials. For example, if an RFP includes a condition requiring audited financial statements for the last three years, the system automatically assembles the relevant file set while simultaneously signaling the approaching expiration of an existing certificate, preventing outdated or invalid confirmations from being included in the disclosure package.

Despite the technological maturity of solutions, the practical scalability of knowledge libraries is constrained by organizational and behavioral factors. Empirical studies of KMS implementation indicate that key sources of resistance include a deficit of stable knowledge-sharing culture and insufficient user training, leading the system to be perceived as an external burden rather than a tool for reducing transaction costs [25, 26]. Measures for overcoming such barriers include institutionalizing accountability for content quality via the introduction of a "knowledge steward" role responsible for taxonomy integrity, metadata discipline, and the currency of master objects [14]. An additional managerial lever is associated with the use of incentive mechanisms, including elements of gamification designed to reward employee contributions to library enrichment and the refinement of attribute labeling [21]. Finally, a critical condition for adoption is embedding the KMS into operational loops: the system should not function as an isolated application, but should be integrated into email clients and reporting tools, reducing "switching costs" and making knowledge governance a natural part of daily workflow [18].

By 2030, the functional role of knowledge libraries will likely pass into the next transformation stage under the influence of agentic AI development and the expansion of asset tokenization practices. Under indicative scenarios, by September 2025 up to 52% of enterprises were already actively using AI agents for autonomous execution of business tasks [27]. Applied to asset management, this implies the possibility of delegating initial counterparty due diligence procedures to an agent, where the matching of provided data to internal limits and risk-management policies is performed based on norms and criteria fixed in the KMS [27].

The projected growth of fund tokenization—forecast to increase eightfold by the end of the decade—creates requirements for transferring part of legal and compliance logic into a digital execution environment. In such a configuration, the knowledge library acquires the role of a source of formalized "legal logic" for smart contracts: investor eligibility rules, jurisdictional constraints, and other normative conditions can be interpreted and applied automatically in a near-real-time mode, ensuring adherence to predefined restrictions without constant manual oversight.

Conclusion

The analysis confirms that building a scalable knowledge library functions as a foundational condition of an asset manager's operational resilience in the context of the "great convergence." A reorientation from fragmented file-based storage toward a formalized hierarchical taxonomy and activity-based metadata simultaneously reduces transaction costs and increases the governability of the evidentiary base required to satisfy strengthened regulatory requirements, including ILPA 2.0 and the SEC Marketing Rule. This transformation shifts knowledge from the status of a weakly structured archive into a controlled corporate resource, where unified rules of classification, versioning, and disclosure reproducibility are maintained.

The implementation of content inheritance mechanisms forms a technological scaling layer that enables the replication of investment strategies across dozens and hundreds of fund configurations without losing control over data integrity and quality. Canonization of master statements and governed local overrides of specific fragments minimize the risk of divergence across parallel materials and reduce the probability of compliance incidents caused by unsynchronized updates of marketing and reporting packages. Quantitative ROI estimates demonstrate that investments in intelligent knowledge management systems can pay back within the first year of operation through DDQ/RFP automation and by lowering the likelihood of financial losses associated with regulatory sanctions.

Over a longer horizon, knowledge libraries evolve into the cognitive core of financial organizations, supporting autonomous AI agents and smart-contract infrastructure. Over the next decade, managers' competitiveness will be determined not only by alpha-

generation effectiveness, but also by the maturity of intellectual-capital governance: the capability to translate heterogeneous information into a structured, verifiable, and readily retrievable asset becomes a system-forming driver of efficiency. The methodological principles and architectural solutions presented may be treated as a practical foundation for a digital-transformation roadmap for IR and compliance functions in contemporary investment firms.

References

1. McKinsey & Company. (2024, September 18). *Beyond the balance sheet: North American asset management 2024*[Report]. Retrieved from: https://www.mckinsey.com.br/industries/financial-services/our-insights/beyond-the-balance-sheet-north-american-asset-management-2024 (date accessed: February 1, 2025).
2. PwC. (2024, November 19). *PwC 2024 Asset & Wealth Management Report* [Press release]. Retrieved from: https://www.pwc.com/gx/en/news-room/press-releases/2024/pwc-2024-asset-and-wealth-management-report.html(date accessed: February 2, 2025).
3. Institutional Limited Partners Association. (2025, February 3). *How we got here: The collaborative effort behind ILPA's new reporting standards.* Retrieved from: https://ilpa.org/news/how-we-got-to-new-reporting-standards/ (date accessed: February 4, 2025).
4. U.S. Securities and Exchange Commission. (2020, December 22). *SEC adopts modernized marketing rule for investment advisers* (Press Release No. 2020-334). Retrieved from: https://www.sec.gov/newsroom/press-releases/2020-334 (date accessed: February 5, 2025).
5. U.S. Securities and Exchange Commission. (2023, November 14). *SEC announces enforcement results for fiscal year 2023* (Press Release No. 2023-234). Retrieved from: https://www.sec.gov/newsroom/press-releases/2023-234 (date accessed: February 6, 2025).
6. U.S. Securities and Exchange Commission. (2024, November 22). *SEC announces enforcement results for fiscal year 2024* (Press Release No. 2024-186). Retrieved from: https://www.sec.gov/newsroom/press-releases/2024-186 (date accessed: February 7, 2025).
7. U.S. Securities and Exchange Commission. (2024, September 30). *Enforcement results for fiscal year 2024* [PDF]. Retrieved from: https://www.sec.gov/files/fy24-enforcement-statistics.pdf (date accessed: February 8, 2025).
8. Oliver Wyman. (2025, January). *10 asset management trends for 2025.* Retrieved from: https://www.oliverwyman.com/our-expertise/insights/2025/jan/asset-management-trends-for-2025.html (date accessed: February 9, 2025).
9. PwC. (2025, February 20). *Asset and wealth management revolution 2024: Unleashing the transformative power of disruptive technology.* Retrieved from: https://www.pwc.ch/en/insights/fs/asset-and-wealth-management-revolution-2024.html (date accessed: February 22, 2025).
10. Troutman Pepper Locke. (2025, March). *The SEC's private fund rules: What advisers need to know* [PDF]. Retrieved from: https://www.troutman.com/wp-content/uploads/2025/03/the-secs-private-fund-rules-what-advisers-need-to-know.pdf (date accessed: April 2, 2025).
11. Institutional Limited Partners Association. (2025, January). *Reporting template guidance* [PDF]. Retrieved from: https://ilpa.org/wp-content/uploads/2025/01/ILPA-Reporting-Template-v.-2.0-Suggested-Guidance.pdf (date accessed: February 10, 2025).
12. Institutional Limited Partners Association. (n.d.). *ILPA reporting template.* Retrieved from: https://ilpa.org/industry-guidance/templates-standards-model-documents/updated-ilpa-templates-hub/ilpa-reporting-template/ (date accessed: February 11, 2025).
13. Dublin Core Metadata Initiative. (2024). *The African Development Bank SANKOFA Program: Business classification that enables disclosure, protects sensitivity and integrates knowledge into the*

business process. Retrieved from: https://www.dublincore.org/conferences/2024/sessions/best-practices/ (date accessed: February 12, 2025).

14. Responsive. (n.d.). *AI-powered DDQ software.* Retrieved from: https://www.responsive.io/solutions/ddq-software(date accessed: February 13, 2025).
15. Adobe. (n.d.). *Adobe Digital Experience Certification Program.* Retrieved from: https://certification.adobe.com/certifications/landing (date accessed: February 14, 2025).
16. CrafterCMS. (n.d.). *CrafterCMS platform features.* Retrieved from: https://craftercms.com/products/platform-features(date accessed: February 15, 2025).
17. Adobe. (2024, April 10). *MSM best practices | Adobe Experience Manager 6.5 documentation.* Retrieved from: https://experienceleague.adobe.com/en/docs/experience-manager-65/content/sites/administering/introduction/msm-best-practices (date accessed: February 16, 2025).
18. Gartner Peer Insights. (n.d.). *RFP response management applications.* Retrieved from: https://www.gartner.com/reviews/market/rfp-response-management-applications (date accessed: February 17, 2025).
19. G2, Inc. (2025). *Enterprise Grid® Report for Proposal | Winter 2025* [PDF]. Retrieved from: https://www.responsive.io/wp-content/uploads/2025/01/G2CR_GR1129_Responsive_Grid_Report_Enterprise_Proposal_Winter_2025_V1.pdf (date accessed: February 18, 2025).
20. Ha-Thuc, V., & Renders, J.-M. (2013, July 9). *Large scale unsupervised hierarchical document categorization using ontological guidance* (U.S. Patent No. 8,484,245). Retrieved from: https://patents.google.com/patent/US8484245B2/en(date accessed: February 19, 2025).
21. Vu, B., Naik, R. G., Nguyen, B. K., Mehraeen, S., & Hemmje, M. (2025). Automated taxonomy construction using large language models: A comparative study of fine-tuning and prompt engineering. *Eng, 6*(11), 283. https://doi.org/10.3390/eng6110283
22. Guida, M., Caniato, F., Moretto, A., & Ronchi, S. (2023). The role of artificial intelligence in the procurement process: State of the art and research agenda. *Journal of Purchasing and Supply Management, 29*(2), 100823. https://doi.org/10.1016/j.pursup.2023.100823
23. Eiff, J. D., Friederich, K., Krüger, M., & Yalcin-Roder, E. (2025). *Investment Funds Statistics Base (IFS-Base): Data report 2025-01 – Metadata version IFS-Base-Data-Doc-v7-1* [Data report]. Deutsche Bundesbank, Research Data and Service Centre. https://doi.org/10.12757/Bbk.IFSBase.09092406. Retrieved from: https://www.bundesbank.de/resource/blob/861344/c8774b415a8d4e074e1a034757642484/472B63F073F071307366337C94F8C870/2025-01-ifsb-data.pdf (date accessed: February 20, 2025).
24. Alternative Investment Management Association. (n.d.). *Due diligence questionnaires.* Retrieved from: https://www.aima.org/resources/education/due-diligence-questionnaires.html (date accessed: February 23, 2025).
25. Grand View Research. (n.d.). *Intelligent document processing market size, share & trends analysis report.* Retrieved from: https://www.grandviewresearch.com/industry-analysis/intelligent-document-processing-market (date accessed: February 24, 2025).
26. Alam, S., Islam, K. A., Miah, R., Afrin, S., Islam, M. A., Hoque, M. R., & Mahmud, I. (2025). Key success factors of knowledge management systems for optimizing organizational performance: A PRISMA-based systematic literature review. *Heliyon, 11*(16), e44097. https://doi.org/10.1016/j.heliyon.2025.e44097
27. McKinsey & Company. (2024, May 30). *The state of AI in early 2024: Gen AI adoption spikes and starts to generate value.* Retrieved from: https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-2024 (date accessed: February 25, 2025)