

L3A AI USAGE WORKING GROUP

AI Usage Report – Luxembourg Alternative Assets Administrators

Corporate Secretarial Activities · KYC / AML Automation · System Integration & Data Management · Skills & Change Management

October 2026

Contributors

This report was produced by the following members of the L3A AI Working Group:

Working Group Members

- ☐ **Doru Ciobanu** - TMF Group
- ☐ **Lindie Fourie** - Hawksford
- ☐ **Andrea Fusco** - Carta
- ☐ **Silvia Gazzoli** - CSC Global
- ☐ **Quentin Gillois** - Arendt Investor Services
- ☐ **Antonella Graziano** - Arendt Investor Services
- ☐ **Solène Griffete** - CBRE IA&R Luxembourg
- ☐ **Brikena Memaj** - Gen II Luxembourg Services Sarl
- ☐ **Carla Santos** - Bolder Group
- ☐ **Oliver Wolf** - Bolder Group

Coordinators

- ☐ **Christophe Lentschat** - Arendt Investor Services
- ☐ **Marta Bawor** – L3A

About L3A

The Luxembourg Alternative Administrators Association (“L3A”) is the leading representative body for fund and corporate administrators of alternative assets in Luxembourg.

Since its founding in 2004, the association has been dedicated to promoting the interests of alternative service providers and contributing to the development of Luxembourg as the investment hub and operations centre of choice for global alternative investment managers.

Our members deliver our mission through their work on our 5 permanent committees, being Operations, Tax, Regulatory, HR and Events & Communications, with output delivered to industry players, regulators and government bodies amongst others, and published through our website and social media channels, in particular LinkedIn.

As the recognised voice of our profession, we engage with authorities through participation in commissions and working groups, and maintain strong ties with government bodies, professional organisations, and Chambers of Commerce.

Our membership is made up of over 60 members, including Full Members, being Luxembourg administration firms predominantly providing services to foreign investment managers of funds and related entities in the alternative management industry and Associate Members, being Luxembourg professional firms operating in the real asset environment and offering services to Full Members and/or their clients.

Together, we form a trusted network committed to excellence.

Introduction

The Luxembourg Alternative Administrators Association (L3A) has long served as a platform for collaboration, knowledge-sharing, and industry advocacy among fund administration service providers in Luxembourg. In January 2025, the topic of artificial intelligence was formally introduced at the L3A Operations Committee, reflecting a growing recognition across the industry that AI was poised to reshape operational processes, compliance workflows, and client service delivery in fund administration.

To better understand how the industry was engaging with AI, L3A launched a flash survey designed to capture insights directly from market participants. The survey, presented by Christophe Lentschat (Arendt Investor Services) in March 2025, sought to answer three fundamental questions: What is the current usage of AI within our industry? What are the main challenges or barriers encountered in the adoption of AI? And how do participants foresee the development of AI in the industry over the next three to five years? Following a call for increased participation during the L3A Annual General Meeting in June 2025, the survey results were published in November 2025 as the *“Survey on AI Usage - Analysis & Insights.”*

The survey findings painted a nuanced picture. On current usage, responses indicated that AI adoption remains concentrated in niche applications rather than being broadly embedded across operations. The most frequently cited use cases included invoice processing and exception monitoring, fund of funds pricing, basic accounting task automation, smart chat and translation tools for client interactions, and AI-assisted board meeting support such as minute drafting. At the same time, several respondents reported limited or no AI usage in their daily operations, suggesting that the industry was still in the early stages of adoption.

On the question of barriers, regulatory and compliance concerns emerged as the most frequently mentioned challenge, with GDPR and data privacy requirements making AI implementation particularly complex. Data security risks, the lack of standardisation across AI models and firms, scepticism about AI reliability for critical tasks, and internal organisational resistance to change were also identified as significant obstacles.

Looking ahead, the survey revealed expectations of significant transformation over the next three to five years. Respondents anticipated that AI would increasingly handle repetitive and administrative tasks, play a larger role in client onboarding and KYC/AML research, and integrate into decision-making processes including boardroom discussions and report drafting. Some foresaw AI-driven workforce efficiency gains of up to 20% within three years, powered by multi-agent AI frameworks and real-time data analysis. A recurring theme was that firms not adopting AI risk falling behind competitively.

These survey results underscored a clear industry need: practical guidance on how to adopt AI responsibly, effectively, and in compliance with the evolving regulatory framework. In response, L3A launched a dedicated AI Working Group in December 2025.

The Working Group was structured around four interconnected topics that directly addressed the opportunities and concerns surfaced by the survey:

- **Corporate Secretarial Activities (“CoSec”)** addressing the regulatory compliance concerns that survey respondents identified as the primary barrier to adoption, with a focus on the EU AI Act framework, GDPR obligations, and governance best practices for corporate secretarial processes.
- **KYC / AML Automation** - responding to the survey’s identification of client onboarding and compliance screening as high-potential areas for AI application, while addressing the data protection and reliability challenges associated with these processes.

- **System Integration & Data Management** - tackling the data security and standardisation concerns highlighted by the survey, including the architectural and regulatory requirements for deploying AI across multi-client fund administration environments.
- **Skills, Training & Change Management** - directly addressing the organisational resistance and trust challenges identified in the survey, as well as the skills development needed to ensure that AI adoption is sustainable, responsible, and aligned with regulatory expectations.

This Consolidated Report brings together the work produced by the four sub-groups of the L3A AI Working Group.

It is intended as a practical reference for L3A members and the broader Luxembourg fund administration industry, offering analysis of the applicable legal and regulatory framework, identification of key points of attention, and actionable best practices for the responsible adoption of AI.

The report reflects the collective expertise of practitioners across the industry and aims to support firms as they navigate the opportunities and obligations presented by AI in a regulated financial services environment.

Disclaimer: This report has been prepared for information and discussion purposes only. It does not constitute, and should not be construed as legal, regulatory, tax, or professional advice of any kind. The content reflects the views and analysis of the L3A AI Usage Working Group contributors at the time of writing and does not represent the official position of L3A, its member firms, or any regulatory authority. Readers should not act or refrain from acting on the basis of this report without seeking independent professional advice tailored to their specific circumstances. Neither L3A nor the contributors accept any liability for any loss or damage arising from the use of, or reliance on, the information contained in this report.

Topic 1: Corporate Secretarial Activities (CoSec)

I. Context

Corporate Secretarial Activities (“CoSec”) is not a standalone use case but a horizontal framework - encompassing AI Governance, GDPR and Data Privacy - applicable to all AI-assisted Corporate Secretarial (“CoSec”) processes. For the scope of this working group, we will analyse **CoSec** processes, including:

- ☐ Board meeting support & minute drafting and shareholder(s)' resolutions
- ☐ Corporate filings & RCS/RESA maintenance
- ☐ Director mandate management
- ☐ Shareholders register maintenance
- ☐ Execution of documentation, such as agreements, legal documentations, minutes, resolutions, POA
- ☐ Coordination of legalisation and apostille

CoSec services involve significant volumes of personal data (director profiles, shareholder registers and proper documentation) and corporate data, making governance and data privacy particularly critical in this context.

Note: The EU AI Act framework established in this topic applies horizontally to all AI-assisted processes covered by Topics 2, 3, and 4 of this report. For entity-type-specific regulatory mapping (DORA vs non-DORA entities), readers are directed to Topic 3.

II. Key Points of Attention

A) Identify Your Role under the EU AI Act

Before deploying any AI tool within a CoSec process, the organisation must first determine its role (AI Act, Article 3) :

- ☐ **Provider:** Is the company developing its own AI model for CoSec tasks (e.g. a proprietary minute-drafting tool) and/or making it available to clients or third parties? Example: Building an in-house AI solution that automatically generates board minutes and offering it as part of a fund administration service offering.

Deployer: Is the company using a third-party AI system to support CoSec operations under its own authority? Example: Using a third-party tool to transcribe and draft minutes, or an AI-powered entity management platform to track RCS filing deadlines and director mandates.

In the CoSec context, most firms will act as Deployers. This does not eliminate accountability - deployers remain responsible for ensuring appropriate use, human oversight, and GDPR compliance in how the AI system is operated.

B) Risk Classification of AI Systems Used in CoSec

CoSec teams must assess the risk level of each AI tool deployed:

- ☐ Prohibited AI (AI Act, Article 5): No relevant use cases anticipated in CoSec, but vigilance is required.
- ☐ High Risk (AI Act, Article 6 and Annex III): AI systems that assist in assessing the propriety of directors, screening individuals against conflict-of-interest registers, or processing personal data for compliance decisions may fall under high-risk categories. These require conformity assessments, logging, and human oversight.
- ☐ Limited Risk (AI Act, Articles 50-52): AI tools used for minute drafting, document summarisation, deadline tracking, or board pack compilation are likely to fall here, but must still comply with transparency obligations - particularly where outputs are shared with clients or regulators.
- ☐ Minimal Risk (overall AI Act, residual category): AI tools used for spam filters and other residual low-risk features.

C) Data Protection Considerations

CoSec processes handle significant volumes of personal data, requiring application of the following measures:

- ☐ **Data minimisation (GDPR, Article 5(1)(c))**: AI tools must only process personal data strictly necessary for the CoSec task at hand (e.g. do not feed full KYC files into a minute-drafting tool).
- ☐ **Purpose limitation (GDPR, Article 5(1)(b))**: Personal data collected for one CoSec purpose (e.g. director onboarding) must not be repurposed for unrelated AI training or processing.
- ☐ **Lawful basis for processing (GDPR, Article 6)**: The organisation must identify and document a valid legal basis for each AI-assisted processing activity involving personal data.
- ☐ **Data transfers (GDPR, Articles 44-49)**: Where AI tools are hosted outside the EEA, appropriate transfer mechanisms (SCCs, adequacy decisions) must be in place before processing director or shareholder data.
- ☐ **Vendor due diligence (GDPR, Article 28)**: **All AI tools must be vetted before being onboarded and, if sufficiently solid, a Data Processing Agreements (DPAs) must be negotiated if they process personal data on behalf of the organisation.**
- ☐ **Accountability & Records of Processing Activities (GDPR, Articles 5(2) and 30)**: AI-assisted CoSec processes involving personal data must be documented in the organisation's Records of Processing Activities (RoPA), reflecting the AI tool as a processing means.
- ☐ **Data Protection Impact Assessment (DPIA) (GDPR, Article 35)**: Where AI tools are used for systematic processing of personal data and might result in risks for data subjects, a DPIA should be conducted prior to deployment.
- ☐ **Data Retention (GDPR, Article 5(1)(e))**: AI-generated outputs (e.g. draft minutes, action logs) containing personal data must be subject to the same retention and deletion policies as manually produced documents.

For the architectural approach to data segregation and purpose limitation compliance in multi-client environments, readers are directed to the Track A/Track B two-track architecture described in Topic 3.

D) AI Literacy and Training for CoSec Staff

(AI Act, Article 4 - AI Literacy obligation applicable to all deployers)

CoSec professionals using AI tools must receive targeted training covering:

- ☐ **Human in the loop:** AI-generated minutes, resolutions, or filing documents must always be reviewed and validated by a qualified CoSec professional before use or submission. No AI output should be filed with the RCS or shared with a board without human sign-off.
- ☐ **Data quality:** High-quality data sources are essential for producing reliable insights. AI systems must be able to access data that is stored in a structured, machine-readable format. Attention must be paid to integration and compatibility with other systems.
- ☐ **Prompting quality:** Where generative AI tools are used (e.g. for drafting board minutes or resolutions), staff must be trained to provide accurate, complete, and unambiguous instructions to obtain reliable outputs.
- ☐ **Hallucination risk:** CoSec staff must be aware that AI tools can generate plausible but factually incorrect content - particularly dangerous in a legal and regulatory context (e.g. incorrect statutory references, wrong filing deadlines).

E) Internal AI Governance and Policies

The organisation should establish and maintain internal policies including ad hoc policies addressing AI use in CoSec, covering at minimum:

- ☐ Approved AI tools and prohibited use cases within CoSec workflows;
- ☐ Mandatory human review checkpoints (e.g. before filing, before board distribution);
- ☐ Data handling protocols when inputting corporate or personal data into AI tools;
- ☐ Incident and error reporting procedures for AI-generated mistakes in corporate records;
- ☐ Audit trail requirements to ensure AI-assisted decisions remain traceable and explainable.

III. Best Practices for AI in CoSec Services

The following best practices aim to help CoSec teams use AI tools efficiently while ensuring compliance with AI governance, data protection, and professional standards. These practices build on the governance principles outlined in Section B and translate them into operational actions for day-to-day CoSec activities.

A) Best Practices for AI Deployment

To ensure effective and responsible use of AI tools in CoSec operations, teams should adopt the following practices:

Define the use case before selecting a tool

Clearly identify the operational objective before deploying AI. For example, an AI tool designed for drafting meeting minutes should not be used for regulatory interpretation or conflict-of-interest screening.

Ensure data quality and accessibility

AI tools rely on structured and reliable data. Corporate records, statutory registers, and governance documentation should therefore be maintained accurate and up to date, in structured and machine-readable formats. Integration with entity management systems and official registries (e.g. RCS) should be prioritised where possible.

Adopt a cross-functional deployment approach

AI implementation should involve CoSec, Legal, Compliance, IT, Security and Data Privacy teams. Data protection specialists should review any tool processing personal data before deployment.

Monitor performance continuously

Regularly assess the reliability of AI outputs by comparing them with validated internal standards (e.g. benchmark minutes or template resolutions). Error trends and anomalies should be tracked and addressed as part of the organisation's AI governance framework.

B) Know Your Role (Provider vs Deployer) - Apply the Appropriate Controls

For each AI use case, CoSec teams must determine their regulatory role under the EU AI Act (as established in Section II.A above) and apply the corresponding controls. The practical implementation steps are set out below.

Document the role

Formally record whether the organisation acts as a Provider or Deployer of the AI system. This classification should be documented in the project documentation and reflected in the Record of Processing Activities (RoPA).

Where acting as Deployer (most CoSec use cases)

Ensure that:

- ☐ a Data Processing Agreement (DPA) is signed with the vendor;
- ☐ appropriate human oversight mechanisms are implemented;
- ☐ transparency obligations towards clients or meeting participants are respected where relevant.

Where acting as Provider (less common)

Engage Legal and Compliance teams early in the process to manage regulatory obligations such as conformity assessments, technical documentation, and post-deployment monitoring.

C) Classify the Risk and Apply Proportionate Safeguards

Before deploying an AI tool, its regulatory risk level should be assessed and documented.

Classify each use case

Register the tool and its use case as Prohibited, High-risk, or Limited risk in the organisation's AI register before operational use.

Apply enhanced controls for high-risk applications

Where AI is used in processes involving evaluation of individuals or critical governance decisions, additional safeguards should be implemented, including:

- ☐ pre-deployment approval,
- ☐ documented testing against validated outputs,
- ☐ logging and traceability,
- ☐ formal senior management validation.

Apply proportionate controls for limited risk uses

For operational tasks such as summarising or drafting assistance, safeguards should focus on transparency, human review, and clear limitations of use (e.g. no automated filings or external reliance without validation).

D) Protect Data by Default - Minimise, Anonymise, Contain

Data protection principles must guide the use of AI tools within CoSec workflows. The default position should always be to process as little personal data as possible, and to apply technical and organisational measures that limit exposure from the outset. This approach reflects the data protection by design and by default obligation under Article 25 GDPR, and should be embedded into the selection, configuration, and day-to-day use of every AI tool deployed in CoSec functions.

Ensure transparency when an AI tool is used to prepare meeting minutes

Participants must be clearly informed whenever an AI-assisted minutes-taking tool is used. This information should be included in the convening notice and reiterated at the start of the meeting by the chairperson or company secretary. For virtual meetings, the notification can also be reinforced through an on-screen message or pop-up. The tool must only be used for generating draft minutes, and any withdrawal of consent must be respected immediately and reflected in the meeting process. Where the legal basis for processing is consent, organisations should ensure that it is freely given, specific, and informed in accordance with Article 7 GDPR, and that withdrawing it carries no negative consequences for the participant. It is also advisable to document the transparency measures taken for each meeting as part of the audit trail, particularly for board or committee meetings where sensitive deliberations may be recorded.

Apply data minimisation

Only the data strictly necessary for the task should be input into AI tools. Sensitive documentation such as full KYC files should never be uploaded unless strictly required and authorised. Before submitting any document or dataset to an AI tool, users should assess whether a redacted, summarised, or anonymised version would be sufficient to achieve the same outcome. This habit should be reinforced through training and, where possible, built into the workflow itself through access controls that limit what can be uploaded.

Anonymise information whenever possible

Replace personal identifiers with roles, initials, or generic descriptions in prompts or training examples where feasible. Anonymisation, where it is genuine and irreversible, takes the data outside the scope of the GDPR entirely, reducing both compliance risk and the potential impact of any security incident. Organisations should consider developing a simple anonymisation checklist or guidance note for staff to follow before engaging AI tools with any personal data information.

Avoid special categories of personal data

Sensitive data (e.g. health information or political opinions) should not be processed unless strictly necessary and legally justified. In CoSec workflows, such data may arise unexpectedly — for example, in the context of board member absences, shareholder resolutions, or employment-related matters discussed at governance level — and staff should be trained to recognise and handle these situations carefully. Where processing of special category data is unavoidable, an explicit legal basis under Article 9 GDPR must be identified, documented, and reviewed by the DPO prior to use.

Control data transfers and hosting locations

Prefer solutions hosted within the EEA. Where data transfers outside the EEA occur, ensure appropriate transfer mechanisms and contractual safeguards are in place. Organisations should verify the hosting location of any AI tool as part of the procurement process and not rely solely on vendor representations without contractual backing. Organisations must also ensure that data transfers are already covered in contracts with clients.

Apply strict retention rules

Where meeting recordings are used to generate draft minutes, they should be deleted once the minutes have been produced, and in any case within a short retention period aligned with internal policies. Retention schedules for AI-generated outputs and any associated input data should be clearly defined, communicated to staff, and technically enforced where possible. The DPO should be involved in setting these schedules to ensure they are consistent with the organisation's broader RoPA and reflect the principle of storage limitation under Article 5(1)(e) GDPR.

E) Build AI Literacy - Checklists and Human Oversight

AI tools should support, not replace, professional judgement in CoSec activities.

Organisations looking to build meaningful AI literacy should consider the following practices: conducting regular, role-tailored training that distinguishes between technical, operational, and governance-level understanding; maintaining up-to-date internal documentation on the AI systems in use, including their intended purpose, risk classification, and known limitations; establishing clear escalation paths so that staff can flag concerns about AI outputs without friction; and embedding AI literacy assessments into onboarding and annual review processes.

Maintain human review at all times

AI-generated documents must always be reviewed and validated by a qualified CoSec professional before distribution or submission (see Section II.D for the full human-in-the-loop obligation).

Promote structured prompt practices

Prompts should be precise and task-focused, avoiding unnecessary confidential or identifying information. It is advisable to apply the GDPR data minimisation principle to AI interactions, asking whether personal data is strictly necessary or whether anonymised alternatives could achieve the same result. Organisations are encouraged to develop internal prompt guidelines for high-risk areas, clearly indicating which data should never appear in a prompt. Finally, staff should be aware that prompts submitted to external AI tools may be retained by the provider, making prompt hygiene a regulatory necessity where no DPIA or appropriate legal basis has been established.

Implement simple quality control checks

Before validating AI-generated documents, CoSec professionals should verify key elements such as:

- ☐ names and roles,
- ☐ meeting dates,
- ☐ statutory references,
- ☐ actions and responsible persons.

F) Govern the Lifecycle - Approve, Monitor, Evidence

AI tools used in CoSec activities should be governed throughout their lifecycle. This means that governance should not be limited to the initial procurement or deployment phase, but should extend continuously to cover updates, changes in use, and eventual decommissioning. A structured lifecycle approach ensures that accountability remains clear and that compliance obligations under both the EU AI Act and the GDPR are met at every stage.

Maintain an approved tools register

Organisations should keep an internal list of authorised AI tools and explicitly prohibited use cases for CoSec functions. This register should include key details such as the tool's provider, intended purpose, risk classification, and the legal basis under which any personal data is processed. It should be reviewed periodically and updated whenever a new tool is introduced or an existing one is materially changed.

Define mandatory validation checkpoints

Human review should be required at key stages, such as before statutory filings, before board circulation, or before regulatory submission. These checkpoints should be formally documented and assigned to a responsible individual, ensuring that AI-generated outputs are never submitted or acted upon without meaningful human oversight. This practice directly supports compliance with the EU AI Act's human oversight requirements and mitigates the risk of consequential errors reaching regulators or stakeholders.

Maintain a clear audit trail

Usage records should document when the tool was used, by whom, and what type of inputs and edits were made. This audit trail is essential not only for internal accountability but also for demonstrating compliance in the event of a regulatory inquiry or data subject request. Records should be retained in accordance with the organisation's data retention policy and be readily accessible to the DPO and compliance function.

Track incidents and corrections

Any errors or inaccuracies generated by AI tools should be recorded and analysed to improve processes and reduce future risks.

Update governance documentation

Changes to AI tools, processing scope, or providers should trigger updates to internal governance documentation, including RoPA, DPIAs (where required), and the organisation's AI governance register. These updates should follow a defined change management process or scope rather than being made on an ad hoc basis, ensuring that documentation accurately reflects the organisation's current AI landscape. The DPO should be consulted as part of this process, particularly where changes introduce new categories of personal data processing or increase the risk profile of an existing tool.

V. Conclusions

AI tools are already reshaping how Corporate Secretarial functions operate — from drafting board minutes and tracking director mandates to managing statutory filings and maintaining shareholder registers. In the financial services sector, where precision, confidentiality, and regulatory accountability are non-negotiable, a misfiled resolution, an inaccurate statutory reference, or an unintended disclosure of personal data can have immediate legal and reputational consequences.

Before deploying any AI tool in a CoSec workflow, know your role, classify the risk, protect the data, and govern the lifecycle. CoSec professionals in financial services firms, fund administrators, and corporate service providers must be equipped to ask the right questions — Is this tool authorised? Has a DPA been signed? Is there a human reviewer assigned before this document reaches the board or the regulator?

Organisations that treat AI governance in CoSec as a tick-box exercise risk non-compliance with both the EU AI Act and the GDPR, as well as exposure to the professional and fiduciary risks that come with errors in corporate records. Those that embed governance into their day-to-day operations — proportionately, practically, and with clear human accountability at every critical checkpoint — will be best placed to harness the benefits AI offers, while maintaining the standards that clients, regulators, and stakeholders rightly expect.

Topic 2: KYC / AML / Onboarding Automation

I. Context - Short Description of AI Use Cases

Note: The EU AI Act horizontal framework established in Topic 1 of this report applies to all AI-assisted KYC and AML processes described in this section. Organisations must determine their role (Provider or Deployer under the AI Act) and apply the appropriate risk classification before deploying any AI tool in the use cases below.

AI-Enabled KYC Data Migration and Gap Analysis

From discussions with prospective clients who are unhappy with the service standards of their current providers, it is clear that one of the main reasons they are reluctant to switch to another Fund Administrator is that they do not want their investors to go through a long and painful KYC remediation process. At the same time, the CSSF expects Fund Administrators providing Registrar & TA services to review the KYC files themselves, rather than simply accepting a reliance letter from the previous provider.

This is where AI can help: if the new Fund Administrator has a proprietary KYC tool or uses a web-based platform developed by a vendor, AI can be leveraged to transfer data from Excel spreadsheets into the new KYC platform without having to do it manually. AI can also highlight any gaps, so that investors are required to provide only the specific information needed.

Enhancing Sanctions, PEP and Adverse Media Screening Processes

The CSSF expects Fund Administrators to have a robust screening process in place to identify natural and legal persons appearing on sanctions and PEP lists or in adverse media. It has become common practice to complement automated screening with manual Google searches to ensure that no negative information is overlooked. This exercise is time-consuming and repetitive and is therefore prone to human error. As a result, it does not need to be performed manually: Fund Administrators can now use AI to build a robust and detailed script that performs this task quickly and efficiently.

II. Key Points of Attention

While AI can significantly streamline KYC remediation during provider transitions, its implementation presents several challenges for Luxembourg Fund Administrators.

- **Data privacy:** The use of AI tools (particularly web-based or vendor-hosted solutions) requires strict controls to ensure compliance with GDPR, including data minimisation, secure data transfers, and clear accountability for data processing.
- **Skill gaps among staff:** Fund Administrators may face skill gaps among staff involved in KYC and screening operations, who need sufficient understanding of AI-driven processes to oversee outputs, investigate exceptions, and exercise appropriate professional judgement.
- **Reliability:** The reliability of AI solutions remains a key concern: outputs must be accurate, explainable, and subject to robust validation and human oversight to ensure that regulatory expectations around KYC and screening reviews, as well as accountability, are fully met.

III. Best Practices and Action Items

To overcome these challenges, Luxembourg Fund Administrators should adopt a structured and well-governed approach to the use of AI in KYC and screening processes.

A) Data Protection

Best Practice

The GDPR data protection obligations relevant to KYC/AML AI tools are as established in Topic 1, Section II.C of this report (data minimisation, purpose limitation, lawful basis, DPA requirements, DPIA, RoPA, and data retention). From a KYC/AML operational perspective, these translate into the following action items. Preference should be given to solutions that allow for local data hosting or clear segregation of personal data.

Possible Action Items

Data Protection Impact Assessments (DPIAs):

- ☐ Management should ensure that a DPIA is performed before the deployment of any AI tool that processes personal data.

Data Minimisation Controls:

- ☐ The Company should ensure that only the minimum personal data necessary is used when interacting with AI tools.
- ☐ Employees should be prohibited from inputting confidential or personal data into AI systems unless explicitly approved.

Third-Party AI Vendor Due Diligence:

- ☐ Prior to onboarding any third-party AI provider, the Company should perform vendor due diligence covering data protection, security, and regulatory compliance.

Contractual Safeguards:

- ☐ Contracts with AI providers should include data protection clauses, confidentiality obligations, security standards, and restrictions on data use.

Security and Oversight:

- ☐ The Company should ensure that AI tools meet internal IT security standards and are subject to appropriate monitoring and access controls.

Data Hosting and Data Segregation:

- ☐ Where possible, the Company should prioritise AI solutions that allow local or EU-based data hosting.

Ongoing Monitoring and Reporting:

- ☐ Management should ensure that the use of AI tools involving personal data is periodically reviewed for compliance with GDPR and internal policies.

B) Staff Training and Awareness

Best Practice

To address skill gaps, Fund Administrators should invest in targeted training for compliance, operations, and oversight functions, ensuring that staff understand how AI tools operate, their limitations, and how to interpret and challenge outputs. Clear procedures should define roles and responsibilities, ensuring that AI supports, rather than replaces, professional judgement and regulatory accountability.

Possible Action Items

Staff Training on AI Tools:

- ☐ The Company should develop and deliver targeted training programmes for staff in compliance, operations, risk, and oversight functions on the use of AI tools.

AI Literacy and Awareness:

- ☐ The Company should ensure that relevant staff are able to understand, interpret, and critically assess outputs generated by AI tools.

Defined Roles and Responsibilities:

- ☐ The Company should document clear roles and responsibilities for the development, deployment, monitoring, and oversight of AI tools.

Human Oversight and Decision-Making:

- ☐ The Company should ensure that AI tools support, but do not replace, human judgement in operational or compliance processes.

Procedures for AI Usage:

- ☐ The Company should establish internal procedures governing when and how AI tools may be used, including any restrictions on their use in regulated or sensitive activities.

Monitoring and Challenge of AI Outputs:

- ☐ Staff responsible for using AI tools should review and challenge AI-generated outputs before they are relied upon for operational or regulatory purposes.

C) Accuracy and Soundness

Best Practice

To ensure reliability, AI solutions should be subject to robust validation, testing, and ongoing monitoring, with documented controls over accuracy, explainability, and exception handling. Human oversight must remain embedded in the process, particularly for high-risk cases, with clear escalation mechanisms and audit trails. By embedding AI within a strong governance and control framework, Fund Administrators can achieve meaningful efficiency gains while continuing to meet CSSF and regulatory expectations.

Possible Action Items

Validation of AI Solutions Prior to Deployment:

- ☐ The Company should ensure that AI tools are subject to appropriate validation and testing before being deployed in production environments.

Ongoing Monitoring of AI Performance:

- ☐ The Company should implement ongoing monitoring mechanisms to assess the performance and reliability of AI tools.

Documented Controls and Testing Procedures:

- ☐ The Company should document validation methodologies, testing procedures, and control frameworks applicable to AI solutions.

Explainability and Transparency:

- ☐ The Company should ensure that AI outputs are sufficiently explainable to allow staff to understand how results are generated and to assess their reliability.

Exception Handling Procedures:

- ☐ The Company should establish procedures to detect, review, and resolve exceptions or anomalies generated by AI systems.

Human Oversight for High-Risk Use Cases:

- ☐ The Company should ensure that human oversight remains embedded in AI-supported processes, particularly for high-risk, sensitive, or regulatory-relevant activities.

Escalation Mechanisms:

- ☐ The Company should establish clear escalation procedures where AI outputs are unclear, inconsistent, or potentially incorrect.

Audit Trails and Record Keeping:

- ☐ The Company should ensure that AI-supported processes generate appropriate audit trails, allowing the Company to trace decisions, inputs, outputs, and human interventions.

Governance and Oversight Framework:

- ☐ Management should ensure that AI tools are embedded within an appropriate governance and control framework, including policies, procedures, and oversight responsibilities.

IV. Conclusions

In practice, Fund Administrators can leverage AI technologies while remaining fully compliant with data protection and regulatory requirements by adopting a "controlled environment" approach grounded in specific legal and regulatory provisions. Under **GDPR, Article 5(1)(b) and (c)**, personal data must be collected for specified purposes and limited to what is necessary (purpose limitation and data minimisation). This means that AI tools should be configured so that investor data is processed strictly for KYC and screening purposes and is not reused by the vendor for model training or product development. In line with **GDPR, Article 28**, where an AI vendor acts as a processor, a robust data processing agreement must clearly prohibit any secondary use of data, define confidentiality and security obligations, and regulate sub-processing. Furthermore, pursuant to **GDPR, Articles 44-49**, any transfer of personal data outside the EEA must be subject to an adequacy decision or appropriate safeguards (e.g. Standard Contractual Clauses), ensuring that data is not stored or accessed in jurisdictions that would conflict with GDPR standards.

From a Luxembourg regulatory perspective, the use of AI solutions - particularly when cloud-based - must also comply with the outsourcing and ICT risk management requirements applicable to the entity type. For entities within the scope of DORA (credit institutions, Chapter 15 ManCos, AIFMs), DORA governs ICT outsourcing directly, while CSSF **Circular 22/806** as amended by **Circular 25/883** applies to business process outsourcing only. For non-DORA entities (specialised PSF, support PSF, Chapter 16-only ManCos), **Circular 22/806** applies in full scope to both ICT and business process outsourcing. In all cases, these requirements mandate clear governance over outsourced ICT services, documented risk assessments, prior due diligence on service providers, defined data localisation and access controls, and the ability for the Fund Administrator and the CSSF to exercise audit and access rights. Administrators must ensure that AI vendors cannot process data beyond documented instructions, that data hosting locations are transparent and contractually agreed, and that exit strategies are in place.

By combining GDPR-compliant contractual safeguards, strict data transfer controls, adherence to CSSF ICT and outsourcing requirements, and the EU AI Act governance framework described in Topic 1 of this report, Fund Administrators can harness AI-driven efficiencies while maintaining full regulatory compliance and accountability as data controllers.

Topic 3: System Integration & Data Management

I. Context

AI-assisted processes in CoSec, domiciliation, and Transfer Agency depend on data drawn from multiple, typically segregated enterprise systems. Fund administration entities in Luxembourg, whether specialised PSF (registrar agents, domiciliary agents under the Article 28-x family of the 1993 Financial Sector Law), support PSF (Article 29-x), management companies performing administration internally under Chapter 15 or Chapter 16 authorisation, or credit institutions providing fund administration as an ancillary service, manage environments in which NAV calculation, fund accounting, transfer agency, AML/KYC, regulatory reporting, and corporate actions each rely on dedicated platforms with distinct data models.

Regardless of licensing category, AI's value is directly proportional to its ability to access, correlate, and reason across data sources. This value takes two structurally distinct forms, each with a different governance and legal basis requirement:

- **Client-scoped tasks:** drafting board resolutions, investor reports, regulatory filings, correspondence - activities where the AI should work only within a single client's data perimeter. Cross-client retrieval here provides no legitimate value and creates professional secrecy and GDPR exposure.
- **Platform-level analytics:** AML anomaly detection across funds, NAV pattern analysis, cross-fund reconciliation, regulatory consistency checks - activities where cross-client or cross-fund data access is the point. These require a different legal basis, a separate architectural track, and explicit DPO and legal counsel involvement.

Conflating these two use cases - allowing the same enterprise AI tool to perform both on a single undifferentiated data lake - is the source of most AI data governance failures in multi-client fund administration. The governance framework must address them separately.

In practice, the data landscape includes:

- **Fund accounting systems** (NAV, positions, pricing, asset valuations, banking data flows)
- **TA systems** (investor data, transactions, tax statuses)
- **AML/KYC platforms** (status, risk scoring, screening outputs)
- **Market data providers** (Bloomberg, stock exchanges)
- **Document management systems** (prospectuses, KIIDs, Article 23 disclosures, client correspondence)
- **Regulatory reporting outputs** (CSSF, AMLA, ESMA filings)
- **Collaboration platforms** (email, shared drives, messaging systems)

B) Retrieval-Based AI and the Training/Retrieval Distinction

The most widely deployed enterprise AI tools, notably Microsoft 365 Copilot, do not train on customer data. The vendor's contractual commitments on this point are explicit. However, this addresses only one dimension of the risk. What these tools do instead is *retrieve* enterprise data at inference time: querying the full scope of content accessible to the prompting user via the Microsoft Graph, including emails, documents, meeting transcripts, and organisational directory data, driven by semantic relevance rather than folder navigation. Critically, LLMs blend retrieved information into coherent outputs without accurate attribution.

A human reviewing an AI-generated board resolution cannot reliably detect that a specific fee percentage, shareholder restriction, or governance clause was drawn from a different client's file - it will read as a plausible, correctly-formatted provision. This is not a deficiency to be managed by careful reading. Instead, it is an inherent characteristic of generative language models. It is why architectural controls must be established *before* the AI is permitted to retrieve, not after it has generated output.

C) Regulatory Framework by Entity Type

GDPR applies directly to all entities - Articles 5(1)(b) and (c) on purpose limitation and data minimisation, Article 25 on data protection by design, Article 28 on processor obligations. Professional secrecy under Article 41 of the 1993 Law applies to all natural and legal persons subject to CSSF prudential supervision. The ICT and outsourcing framework diverges by DORA scope:

- **DORA entities** (credit institutions, Chapter 15 ManCos, AIFMs): DORA applies directly. New **Circular CSSF 25/882** supplements DORA with Luxembourg-specific requirements on ICT third-party service notifications and the information register. **Circular 22/806** (as amended by **Circular 25/883**) applies for business process outsourcing only; ICT outsourcing is governed by DORA itself.
- **Non-DORA entities** (specialised PSF, support PSF, Chapter 16-only ManCos): subject to **Circular 20/750** (as amended by **Circular 25/881**) for ICT and security risk management, and to **Circular 22/806** (as amended by **Circular 25/883**) in full scope for both ICT outsourcing and business process outsourcing. **Circular 24/847** on ICT incident reporting also remains applicable.

Internal governance requirements likewise differ: **Circular 12/552** applies to credit institutions (and partly to professionals performing lending operations under Article 28-4 LFS); **Circular 04/155** (as amended by **Circular 22/806**) applies to specialised and support PSF; **Circular 18/698** governs ManCos and AIFMs. Applying the wrong circular creates both a compliance gap and a false sense of assurance.

II. Key Points of Attention

The following points apply to any CSSF-supervised entity operating a multi-client fund administration model, with the applicable regulatory instrument varying by licence type as set out in Section I.

A) Over-Permissioned Access at AI Scale

Enterprise environments accumulate permissions reflecting role history and organisational change rather than current need-to-know. In human workflows, latent permissions are practically inert. When AI retrieval tools are deployed, every latent permission becomes active on every query, converting dormant over-permissions into active retrieval pathways across the entire historical data estate. This risk is identical whether the entity is a specialised PSF with 200 fund clients or a ManCo administering 30 sub-funds: it is the multi-client data estate, not the licence type, that determines the exposure.

B) Cross-Client Data Surfacing, LLM Blending, and the Article 41 Distinction

Semantic retrieval does not return documents - it synthesises content. If an AI tool draws on Fund B's fee schedule, governance provisions, or shareholder restrictions while drafting a Fund A resolution, these elements will be integrated into a coherent document without attribution. A human reviewer checking for procedural accuracy cannot reliably detect what source material the LLM used unless the system-level audit log is consulted. The assurance that the output is clean cannot come from reading the output; it must come from confirming that the AI could only access Fund A's data in the first place.

Article 41 of the 1993 Law imposes professional secrecy on all persons working for CSSF-supervised entities, criminally sanctioned under Article 458. The obligation applies regardless of intent. A critical distinction must be maintained:

- **A governance failure caught internally:** AI retrieves Fund B's data during a Fund A task; the system-level audit log identifies cross-client access; no affected content is delivered externally; the matter is remediated. This is an *operational risk event*, logged as a process failure, investigated, and used to improve controls. Internal identification and remediation is a control function, not a disclosure. Article 41 is not engaged.
- **An actual secrecy breach:** Fund B's confidential data is incorporated into an output delivered to Fund A's directors. This *may* engage Article 41 depending on content and circumstances. It requires immediate legal counsel involvement, privilege considerations for associated documentation, classification assessment before any regulatory notification, and a distinction between the factual notification (if required) and the internal legal analysis (which may be privileged).

On incident classification and regulatory notification: where a DORA entity faces a potential major ICT-related incident, the 4-hour notification clock under **Implementing Regulation (EU) 2025/302** runs from *classification*, not from discovery. Legal counsel can be involved in parallel with initial triage. The classification decision is itself a legal and technical assessment. The factual notification to the CSSF and the internal legal assessment of criminal exposure are separate documents serving separate purposes; obtaining legal advice on the latter before notifying the former does not constitute prioritising self-protection over regulatory transparency.

C) GDPR Purpose Limitation: The Two-Track Architecture

GDPR Article 5(1)(b) requires processing only for the purpose for which data was collected. For a multi-client fund administrator, data collected for "administering Fund A's SICAV" has a defined processing purpose. If AI retrieves Fund A's data in response to a query about Fund B, this constitutes further processing for an incompatible purpose.

The resolution is architectural, not prohibitive, but it requires acknowledging that two categories of AI use demand different legal bases and different technical designs:

- **Track A - Client-scoped tasks** (board resolutions, investor reports, client filings): AI retrieval is restricted to the relevant client's data silo by technically enforced controls. No cross-client retrieval occurs. GDPR purpose limitation is satisfied because data collected for Fund A is processed only for Fund A's tasks. The legal basis is the performance of the administration contract (Article 6(1)(b)) or the administrator's legitimate interests in efficient service delivery. This track is fully compatible with the multi-client model provided retrieval scope controls are architecturally implemented and technically enforced before deployment. The segregation required is achievable within a single enterprise AI environment using platform-native controls; it does not require separate tenants or separate AI product subscriptions. What matters is that the scope boundary is technically enforced at the platform layer, not merely described in user instructions or policy documents, so that the AI cannot retrieve outside the permitted perimeter regardless of how a query is framed.
- **Track B - Platform-level analytics** (AML anomaly detection across funds, NAV pattern validation, cross-fund reconciliation, KYC documentation consistency checks): these use cases require processing data from multiple client engagements. A separate legal basis is required - typically Article 6(1)(f) legitimate interests (the administrator's obligation to detect financial crime and operational risk, which serves the interests of all clients) combined with, where possible, anonymisation or aggregation at the output layer so that individual client data is not surfaced in the

result. This track requires a separate, purpose-specific data pipeline with its own technically enforced access scope, a legitimate interest assessment documented by the DPO, and contractual provisions in the administration agreements with each client. It cannot be delivered by the same enterprise AI tool operating on the same undifferentiated data lake as Track A; but it can coexist within the same broader technology environment provided the pipeline has its own technically enforced access scope and its outputs are purpose-specific.

The critical point is that **Track B analytics are achievable, but they require deliberate legal basis mapping and architectural separation from Track A**. Entities that deploy a single enterprise AI tool across a unified data lake, expecting it to serve both tracks simultaneously, may not be operating within a compliant framework regardless of how the results are labelled. Entities that design two-track architectures with appropriate legal basis documentation are providing both the efficiency value of cross-client analytics and a defensible GDPR position.

D) Audit Trail, Logging, and Proportionate Detection Infrastructure

A critical governance question is whether the entity can reconstruct what data the AI retrieved for any given work product. Under standard M365 E3 licensing, Copilot interactions generate metadata logs (user ID, timestamp, AccessedResources file list) but not content logs (actual prompts and responses, which require eDiscovery). Standard retention is 180 days.

Detection infrastructure requirements should be **calibrated to the entity's AI deployment scope and volume**, not applied uniformly. For entities deploying AI at scale across client-facing work, robust logging - including content-level capture through eDiscovery configuration or E5-level tooling - is necessary because the volume of AI-assisted outputs makes manual controls insufficient. For entities with limited AI use, periodic manual review of the E3-level AccessedResources log and operational controls (documented retrieval scope configuration, sampling-based output review) are proportionate alternatives.

Record-keeping obligations derive from different instruments by entity type: **Circular 12/552** for credit institutions; **Circular 04/155** for specialised and support PSF; **Circular 18/698** for ManCos and AIFMs.

E) Data Lake Governance and the Two-Track Design Requirement

Where organisations build or operate data lakes, governance requirements are substantial and must reflect the two-track architecture described in Section II.C. A single undifferentiated data lake feeding a single AI retrieval layer cannot simultaneously support both Track A (client-scoped, GDPR compliant through segregation) and Track B (cross-client analytics and KYC requiring separate legal basis). The architecture must therefore provide:

- ❑ **Segregated client partitions** for Track A: data classified by client, with retrieval scope controls that enforce partition boundaries at the AI query layer.
- ❑ **A dedicated analytics pipeline** for Track B: purpose-specific, with pseudo-anonymisation or aggregation at the output layer, separate access controls, documented legal basis, and DPO sign-off.
- ❑ **Data lineage and classification:** sensitivity labels technically enforced; every AI-generated output carries metadata identifying source data and the track under which it was produced.
- ❑ **Retention and deletion policies:** aligned with GDPR Article 5(1)(e) and the applicable governance circular, with automated purging.

F) ICT Third-Party Risk and the DORA/Non-DORA Divide

Microsoft's designation as a critical ICT third-party provider under DORA (November 2025) affects the full licensing spectrum. For DORA entities: DORA Chapter V applies directly; **Circular 25/882** specifies notification and register requirements; Article 30 contractual obligations apply. Standard M365 enterprise licensing does not meet DORA Article 30 requirements without additional documentation. For non-DORA entities: ICT third-party obligations derive from **Circular 22/806** as amended by **Circular 25/883**; Circular 25/882 does not apply.

An AI-mediated data exposure event where retrieval surfaces confidential client data across client or fund boundaries may constitute an ICT-related incident. For DORA entities, the event must be classified before the notification clock starts; classification is a legal and technical assessment requiring counsel involvement. For non-DORA entities, **Circular 24/847** governs incident reporting.

G) The ManCo-PSF Delegation Boundary: Proportionality and Industry Standardisation

Circular 18/698 requires ManCos to maintain effective oversight of delegated functions. When AI tools are deployed at a delegate PSF, the ManCo's oversight obligation extends to the PSF's data governance and ICT practices.

In practice, ManCo legal and risk departments will tend to impose contractual ICT and AI governance requirements on delegate PSFs that reflect the ManCo's own DORA obligations, regardless of the PSF's non-DORA status, because the ManCo's DORA compliance requires it to demonstrate effective third-party oversight. The more effective response is to work toward **industry-standard baseline clauses for ManCo-PSF AI and ICT governance**, perhaps developed at L3A level. Such a framework would:

- ☐ Define the minimum information rights and audit access a ManCo requires to discharge its DORA oversight obligation regarding a delegate's AI deployment.
- ☐ Specify ICT and data governance standards that satisfy ManCo oversight requirements while being calibrated to the PSF's own circular framework (Circulars 20/750 and 22/806 as amended), not to DORA Chapter V directly.
- ☐ Establish proportionate incident reporting obligations for the PSF toward the ManCo, aligned with Circular 24/847 for non-DORA entities, so that ManCos receive the information they need without requiring PSFs to operate as if DORA-bound.
- ☐ Provide a liability allocation framework that reflects the PSF's non-DORA regulatory perimeter, preventing contractual imposition of liability standards higher than the PSF's licence requires.

A shared industry template produces better outcomes than bespoke contracts in which each ManCo imposes its most restrictive standards and each PSF negotiates from a weaker position. L3A is well-placed to develop this through its working group structure, in dialogue with the CSSF on the proportionality requirements of the outsourcing framework.

III. Best Practices for System Integration & Data Management

The following best practices apply to any CSSF-supervised fund administration entity, adapted to its specific licensing category, DORA/non-DORA status, and operating model.

A) Conduct Pre-Deployment Permission Remediation

Before deploying any retrieval-based AI tool in production, the entity should complete a comprehensive permission audit covering SharePoint site access (identifying sites accessible to broad security groups), shared mailbox permissions, Teams channel membership relative to current role requirements, and OneDrive sharing configurations. Targeted remediation of overly broad permissions must precede deployment, not follow it. For ManCos with internalised administration, the audit should extend to any shared infrastructure between portfolio management and administration functions.

B) Implement Track A/Track B Architectural Separation through Technically Enforced Controls before Deployment

The entity should design its AI data architecture around the two-track distinction described in Section II.C before deploying any AI tool that accesses enterprise data.

Track A (client-scoped tasks) requires technically enforced retrieval scope controls that restrict AI access to the relevant client's data partition. The critical distinction is between *technically enforced controls* and *organisational controls*. Technically enforced controls are platform-native barriers that the AI cannot override regardless of user action or instruction. In M365 environments, these include Restricted SharePoint Search, which limits Copilot's retrieval index to designated sites only, and Restricted Content Discovery, which prevents semantic indexing of content outside the permitted scope. Organisational controls (folder structures, naming conventions, user-facing policies, instructions about which data to query) depend on human compliance at query time and are insufficient for this purpose: they may reduce the risk of inadvertent over-querying, but they do not constitute architectural segregation for the purposes of GDPR Article 25 or professional secrecy compliance.

Track B (platform-level analytics) requires a separate, purpose-specific data pipeline with its own technically enforced access scope, documented legal basis, anonymisation at the output layer where possible, and DPO sign-off. This pipeline can coexist within the same Azure or M365 environment as Track A. The separation required is a technically enforced logical boundary - a dedicated data ingestion layer with its own access controls and query scope, such as a dedicated Azure AI Search index with scoped connectors or a purpose-built analytics environment with separate RBAC-governed data access.

The entity should document which AI tools are authorised for which track, the specific technical controls implementing each track's scope boundary, and the governance process for verifying those controls remain correctly configured following system updates or vendor platform changes.

C) Establish a Governed Data Lake Architecture

Where data lakes are deployed, organisations should implement the Track A and Track B controls set out in Section II.E, including segregated client partitions, a dedicated analytics pipeline, data lineage, and aligned retention and deletion policies.

D) Apply a Two-Layer Assurance Framework to AI-Generated Outputs

LLMs blend retrieved information without attribution. A human reviewer reading an AI-generated output cannot reliably determine whether a specific provision was drawn from the correct client's documents or

from another client's files. This is not a deficiency addressable by more careful reading; it is an inherent characteristic of generative language models. These are two distinct risk categories - data provenance and output quality - requiring two distinct and non-substitutable assurance mechanisms:

- **Layer 1 - Retrieval scope verification (AccessedResources audit log):** the system-level log records which files the AI retrieved during the query. Consulting this log before delivery of a high-risk output confirms that the AI accessed only files within the relevant client's configured retrieval scope. This is the mechanism for inter-client professional secrecy compliance: it establishes that Fund B's data was not retrieved for Fund A's task. Layer 1 does *not* confirm that the output is substantively accurate, legally correct, or free from hallucination. It addresses data provenance only; it does not address output quality.
- **Layer 2 - Substantive output review (human expert):** expert assessment of the AI-generated content for accuracy, legal correctness, completeness, consistency with the client's known facts, and absence of hallucinated provisions. This review is necessary regardless of what Layer 1 shows. A clean access log enables the reviewer to concentrate on substantive quality rather than data provenance. It does not reduce the scope or depth of expert review required.

Neither layer substitutes for the other. The elaborateness of each scales to the output's risk profile:

- **Tier 1 - Full two-layer review:** regulatory filings, client-delivered board resolutions, AML/KYC determinations with legal or regulatory consequences.
- **Tier 2 - Sampling-based review:** routine investor communications, standardised fund reports produced at volume (e.g. 10% random selection plus any output where the AI system flags multi-source retrieval).
- **Tier 3 - User attestation:** internal drafts, meeting notes, preliminary analysis. The user confirms at completion that the AI tool was used within the configured retrieval scope and that the output has been reviewed for substantive accuracy before any external use.

E) Configure Proportionate Logging Infrastructure

Logging infrastructure should be calibrated to the entity's AI deployment scope and volume. For entities with limited AI use, periodic manual review of the E3-level AccessedResources log and operational controls (documented retrieval scope configuration, Tier 2 sampling) are proportionate and should be documented as such. For entities deploying AI at scale across client-facing work, content-level logging through eDiscovery configuration or E5-level tooling provides the reconstruction capability necessary where manual sampling is insufficient. Entities should document their logging approach and the rationale for proportionality.

F) Adopt Cross-Functional Deployment Governance with Circular-Correct References

AI deployment should be subject to the entity's applicable product and process approval framework: sub-chapter 7.3 of **Circular 12/552** for credit institutions; **Circular 04/155** for PSF; **Circular 18/698** for ManCos. A formal AI governance group should maintain an AI register documenting each tool, its data sources, its track classification (Track A, Track B, or both), its EU AI Act risk classification, and its contractual status under either **Circular 25/882** (DORA entities) or the applicable circular framework (non-DORA entities). System integration decisions should involve IT, Compliance, Risk, Legal, the DPO, and the relevant business lines. Track B deployments additionally require documented legitimate interest assessments and DPO sign-off before go-live.

G) Maintain an AI Governance Failure Category with Legally Coherent Framing

The risk function should maintain a discrete operational risk event category for AI governance failures: instances where the system-level audit log reveals cross-client data access that was not authorised by the retrieval scope configuration, or where AI synthesis incorporated content from outside the permitted perimeter. The following framing is essential:

- **Scope:** governance events identified *before* any affected content was delivered externally - process failures and control signals. These are operational risk events, logged and used to improve controls.
- **What this does not cover:** events where cross-client data may have been delivered externally. These require immediate escalation to legal counsel for classification assessment under legal privilege before any documentation is prepared. Factual notification (if classified as a major incident) and internal legal assessment are separate documents.
- **Safe-reporting culture:** staff who identify AI retrieval anomalies and report them internally are performing a control function. Internal reporting of a near-miss does not constitute an admission of wrongdoing under Article 41; Article 41's criminal provision targets actual external disclosure, not internal identification of a process failure.
- **Interpreting incident counts:** zero entries from an entity with active retrieval scope controls, Tier 1 two-layer review, and a documented sampling programme is a good outcome. Zero entries from an entity with no logging, no retrieval scope controls, and no review framework indicates that detection infrastructure does not exist.

Practical Illustration: Microsoft 365 Copilot in a Multi-Client Fund Administration Entity

The following scenario illustrates the two-track architecture in practice, using a third-party fund administrator holding specialised PSF status (a non-DORA entity). The entity provides fund administration, corporate secretarial, and depositary services to 80 fund clients. It has deployed M365 Copilot within a single M365 tenant and built a data lake aggregating fund accounting, TA, and AML/KYC data across all clients.

Track A deployment (client-scoped tasks). A corporate secretary uses Copilot to draft a board resolution for Fund A. Restricted SharePoint Search has been configured to limit Copilot's retrieval index to Fund A's client-scoped SharePoint site only - a platform-native technically enforced control within the single tenant. Before delivery, the reviewer completes both assurance layers. Layer 1: the AccessedResources log is consulted and confirms that all retrieved files are within Fund A's configured retrieval scope - no cross-client files appear. Layer 2: the reviewer assesses the draft resolution for procedural accuracy, legal correctness, and consistency with Fund A's known governance facts. Both layers are documented. The resolution is delivered.

Track B deployment (platform analytics). The entity's compliance function uses a dedicated analytics pipeline - a purpose-built environment within the same Azure tenant, with its own RBAC-governed data access and a separate query scope that cannot access the Track A client-scoped SharePoint sites. The DPO has documented a legitimate interest assessment under GDPR Article 6(1)(f). The pipeline outputs alerts (fund-level anomaly flags) without surfacing individual investor data to the analysts. The administration agreements with each client include a provision authorising cross-fund analytics for AML and risk management purposes.

A governance failure scenario. During a high-volume period, a configuration error temporarily removes the retrieval scope restriction for Fund A's SharePoint site, allowing Copilot to query across multiple client sites. A Layer 1 log review before delivery of a board resolution identifies that two files from Fund B's site appear in the AccessedResources log. The resolution is not delivered. The matter is logged as a governance failure (a process control event - no content was delivered externally, Article 41 is not engaged). The configuration is corrected, the logs are retained, and the incident is used to add a pre-delivery configuration verification step to the workflow. Legal counsel is informed as a precautionary measure but determines that no notification obligation is triggered.

IV. Conclusion

System integration and data management represent the infrastructure layer on which all other AI governance requirements depend. The value of AI in CoSec, domiciliation, and Transfer Agency is unlocked when AI tools are connected to enterprise data at scale. Yet this interconnection generates consequential risks that differ fundamentally between client-scoped tasks and platform-level analytics. Treating these as a single design problem produces governance frameworks that are either too restrictive to preserve AI's value or too permissive to satisfy legal obligations.

Five principles organise the governance response. *First, two-track architecture is the structural solution:* client-scoped tasks under technical segregation (satisfying GDPR Article 25 through architectural control) and platform analytics under separate legal basis, separate pipeline, and DPO-documented legitimate interest assessment. *Second, assurance requires two irreducible layers:* Layer 1 (retrieval scope verification) addresses cross-client professional secrecy; Layer 2 (human expert review) addresses substantive output quality. *Third, governance events and Article 41 exposure are categorically distinct:* internal identification of a cross-client retrieval anomaly before external delivery is a control function, not a disclosure. *Fourth, detection infrastructure and review standards must be proportionate to AI deployment scope.* *Fifth, the ManCo-PSF governance divide is better resolved by industry standardisation than individual contract negotiation:* L3A is positioned to develop template clauses that satisfy ManCo DORA oversight obligations while calibrating PSF requirements to the PSF's own circular framework.

AI supports human decision-making and enhances operational capability when the data infrastructure it relies on is governed with the same analytical rigor applied to the decisions themselves - built on the correct regulatory baseline, designed for the correct use-case track, and calibrated to the entity's actual AI deployment profile.

Topic 4: Skills, Training & Change Management

I. Context

AI is starting to fundamentally reshape operational and compliance-driven functions such as **CoSec/domiciliation and Transfer Agency**. These activities depend on the accurate extraction, interpretation, and processing of documents and investor data. AI introduces powerful opportunities to automate repetitive tasks - drafting minutes or resolutions, classifying legal documents, extracting investor identification data, or preparing capital call notices. Yet, as these tools become embedded in day-to-day workflows, they alter not only the way tasks are completed but also the skills required to perform them.

A defining element of the current landscape is the **speed at which AI technologies evolve**, which is far faster than the traditional rhythm of corporate learning cycles. Employees are asked to work with tools that change monthly or even weekly. This produces a clear divide: some individuals quickly incorporate new functionalities and adjust their approach, while others struggle to follow the changes. In CoSec and TA, where accuracy and regulatory compliance are paramount, this uneven readiness can slow adoption and even increase risk, particularly when employees misinterpret AI outputs or over-rely on them.

Note: The EU AI Act horizontal governance framework described in Topic 1 of this report - including the AI Literacy obligation under Article 4, data protection obligations, and human oversight requirements - applies fully to the context covered in this topic. Readers are directed to Topics 1, 2 and 3 for the detailed regulatory and governance framework; this topic focuses on the practical skills, training, and change management dimension.

A) Changing Nature of Work and New Skill Sets

The changing nature of work demands **new skill sets**. Employees who previously performed manual drafting or manual reconciliation now find themselves supervising AI-generated content, evaluating whether it aligns with legal or policy requirements, and applying professional judgement to confirm its validity. The requirement is no longer simply "knowing the process," but also being able to **critically assess** automated results. Staff must understand the strengths and weaknesses of AI, recognise when an output looks plausible but is inaccurate, and apply the right escalation steps.

B) Training as a Strategic Requirement

Training therefore becomes a strategic requirement. Generic AI training programmes - often designed at an enterprise level - rarely match the specific needs of CoSec or TA activities. Employees need guidance grounded in real operational scenarios: evaluating a draft board resolution, reviewing AI-proposed amendments, checking the completeness of KYC extraction, or assessing the quality of investor communications produced by automated systems. This training must also cover responsible use: confidentiality when entering prompts, avoidance of unnecessary personal data, and recognition of model limitations.

C) Human Dimension of Change

At the same time, the human dimension of change cannot be underestimated. AI naturally generates **uncertainty and emotional resistance**. Some employees worry about job security or fear that automation will diminish their expertise. Others distrust AI output and hesitate to incorporate it into regulated workflows. Without clear communication and strong sponsorship, these concerns can reduce engagement and adoption.

Employees must be aware and understand how AI will affect their work, why it is being introduced, and what responsibilities remain with them - especially regarding oversight and final validation. Awareness sessions are therefore a key element.

D) Governance

Governance forms another essential piece of the context. AI in regulated environments requires clear rules: who validates outputs, who owns the final decision, how evidence is recorded for audit, and what control framework applies. Ambiguity around these responsibilities can lead to inconsistent practices and increase regulatory exposure. **Clear ownership roles - such as AI Product Owners, Data Stewards, or Model Validators - must be defined early.**

II. Key Points of Attention

Several areas require focused attention to ensure that AI adoption in CoSec and TA is safe, consistent, and sustainable.

A) Skills Development

Employees must learn to work effectively at the intersection of automation and human judgement. This includes understanding how AI tools generate outputs, how to interpret their strengths and limitations, and how to identify errors or inconsistencies that require intervention. In CoSec, this might involve spotting missing legal clauses or incorrect corporate references in AI-drafted documents. In TA, it can involve reviewing data extracted from investor files, checking for completeness, identifying anomalies in risk assessments, or identifying issues in reconciliation performed by AI between investor notices and bank statements.

B) Role-Specific, Scenario-Based Training

Training must address the actual activities employees perform. For CoSec teams, this means training built on real governance documents: resolutions, minutes, statutory filings, or correspondence. For TA teams, training must simulate actual onboarding files, AML/KYC situations, investor communications, and capital call workflows. Training must emphasise not only how to use the tool but how to use it safely, including avoiding the misuse of personal data and recognising the boundaries of what AI can do.

C) Human-in-the-Loop Controls

In regulated functions, **human oversight remains mandatory.** Employees must be trained to review **AI-generated content with discipline:** verifying source information, documenting decisions, recording any overrides, and escalating issues when necessary. In many tasks - such as validating investor names, checking KYC completeness, or finalising board documentation - the quality of human review is the primary safeguard against errors. Training must therefore develop strong analytical skills and encourage a mindset of constructive skepticism.

Governance clarity represents another important point of attention. Staff must know exactly who is responsible for validating AI outputs, who approves final documents, and how the audit trail should be maintained. Otherwise, inconsistencies arise: one team may trust AI too much, another may avoid using it entirely, and overall risks increase.

D) Awareness

Awareness is not only about knowing that AI tools exist; it is about understanding **how**, **when**, and **why** to use them responsibly. In both CoSec and TA, a lack of awareness can introduce risks even when skills and training are formally in place. Employees must be aware of:

- **The limits of AI:** that outputs may be incomplete, approximate, or context-dependent.
- **Their continuing responsibility:** the human remains accountable for the final result, even when AI is involved.
- **Confidentiality obligations:** which client or investor information can be safely entered into AI tools and which cannot.
- **Misinterpretation risks:** awareness that AI may produce plausible but incorrect suggestions.
- **Escalation pathways:** when results look unusual, incomplete, or inconsistent.

Awareness also includes a cultural dimension. Employees must understand that AI is introduced to support their roles, not replace them. In functions such as CoSec and TA - where expertise, judgement, and attention to nuance are essential - AI is an assistant, not a substitute. Raising awareness therefore requires deliberate communication: explaining in simple terms what AI can do, what it cannot do, how it is governed, and what the expected behaviors are. This is where townhalls, Q&A sessions, internal success stories, and "AI champions" become essential to stabilising adoption.

E) Change Management

Employees must clearly understand why AI is being introduced and how it benefits their role, their team, and the organisation. Transparent communication helps dispel fears that AI will replace them and reinforces the message that human expertise remains essential. Early involvement of staff in pilot programmes, confidence-building demonstrations, and visible leadership support all contribute to stronger adoption. Creating AI "ambassadors" or super-users helps reinforce learning and build trust across teams.

III. Best Practices for AI in CoSec Services/Domiciliation/TA

A) Establish a Formal AI Skills & Change Management Framework

To overcome challenges in CoSec services, domiciliation, and TA, Professionals of the Financial Sector should establish a formal AI Skills & Change Management framework anchored in governance and capability development.

Create an AI governance group that concentrates on people, processes, and technology, with representation from risk, compliance, HR, IT, and business lines. This group will define learning objectives, certification paths, and role-specific competency models, ensuring that training aligns with regulatory expectations, risk appetite, and operational needs.

By codifying governance and competencies, the organisation can steadily build a workforce equipped to design, deploy, and oversee AI applications with clarity on ownership, accountability, and expected outcomes. The framework should regularise onboarding, upskilling, and refreshers to keep pace with evolving AI capabilities and evolving regulatory guidance.

B) Human-in-the-Loop Design and Clear Decision Rights

Design processes where AI recommendations are clearly labelled, with explicit points for human approval or override, and document escalation criteria so front-line professionals know exactly when to rely on AI versus when to escalate to a human expert. This approach preserves professional judgement and accountability, while enabling scalable automation.

To operationalise this, map decision points, annotate confidence levels, and implement auditable trails that capture rationale for both AI-driven and human decisions. Regular reviews should assess where AI adds value, where it should defer, and how to handle exceptions or disputes to minimise risk and maintain client trust.

C) Targeted, Role-Based Training Programmes

Develop training programmes that are tailored to each role's actual responsibilities.

For frontline staff, provide instruction on how to read AI outputs, interpret confidence scores, and escalate appropriately when AI signals uncertainty or potential risk.

For risk and compliance teams, offer modules on model risk management fundamentals, bias and fairness checks, and awareness of regulatory implications, with practical guidance on documentation and ongoing monitoring.

For data and technology teams, emphasise model lifecycle management, data governance, privacy-by-design, and robust security controls, including incident response and privacy protections.

Enrich these programmes with practical simulations, scenario-based exercises, and hands-on labs to reinforce correct usage, validation, and oversight, ensuring learners can apply concepts effectively in daily tasks.

D) Ethics and Client Transparency

Anchor AI usage in a strong ethics and transparency framework. Develop explicit ethical guidelines for AI use, including fairness, non-discrimination, explainability, and accountability, and disclose AI-assisted processes to clients where appropriate.

Provide user-friendly explanations of AI-driven actions to clients and internal stakeholders, so decisions are understandable and interpretable. Align ethical guidelines with professional standards, regulatory expectations, and client needs in CoSec services, domiciliation, and TA. This transparency not only supports informed decision-making but also strengthens trust, supports auditability, and provides a clear reference for addressing client questions about AI-enabled outcomes.

IV. Conclusion

AI is redefining how CoSec/domiciliation and Transfer Agency operate - not only by enabling automated, high-precision processing of documents and investor data, but also by redefining the skills, governance, and ethical foundations required to do this safely and compliantly. AI accelerates work and enhances accuracy, but it also demands new skills, targeted training, and clear governance to prevent misinterpretation and regulatory risk. By recognising the rapid evolution of AI tools and the need for role-specific competencies, organisations can align technology with regulatory expectations and operational realities.

This alignment requires a four-part approach: (1) **governance and skills development** to establish clear ownership, accountability, and competencies; (2) **human-in-the-loop design** that labels AI outputs, defines decision rights, and preserves professional judgement; (3) **targeted, role-based training** grounded in real CoSec/TA scenarios, including risk management, KYC, governance documents, and capital call workflows; and (4) **ethics and transparency** to build client trust and auditability, including explainability, fairness, and clear disclosure of AI-assisted processes. Together, these elements enable safe, scalable adoption that enhances accuracy, strengthens compliance, and supports staff confidence as AI capabilities evolve.

Overall Conclusion

The work of the L3A AI Working Group, across its four interconnected topics, confirms a central finding: AI presents a transformative opportunity for the Luxembourg fund administration industry, but realising that opportunity requires deliberate, structured, and legally grounded action. The industry stands at an inflection point where early and responsible adoption can deliver substantial competitive advantage, while inaction or uncontrolled deployment carries material regulatory and operational risk.

Several cross-cutting themes emerge from the consolidated analysis.

Governance is the prerequisite, not the afterthought. Across all four topics, the report demonstrates that effective AI adoption begins with governance. The EU AI Act, GDPR, CSSF circulars, and DORA collectively establish a comprehensive regulatory framework that requires organisations to classify their AI systems, determine their regulatory role, enforce data protection by design, and maintain appropriate human oversight. Whether the use case involves CoSec minute drafting, KYC data migration, cross-client analytics, or employee training, the governance framework must be established before deployment - not retrofitted afterwards. The horizontal applicability of the EU AI Act, as established in Topic 1, reinforces that no AI-assisted process is exempt from these obligations.

Data protection and professional secrecy demand architectural solutions. The report consistently highlights that policy-level commitments alone are insufficient. GDPR purpose limitation, data minimization, and the professional secrecy obligations under Article 41 of the 1993 Financial Sector Law require technically enforced controls - particularly in multi-client fund administration environments where retrieval-based AI tools can surface data across client boundaries. The two-track architecture described in Topic 3, distinguishing between client-scoped tasks (Track A) and platform-level analytics (Track B), provides the structural framework for achieving compliance while preserving the operational value of AI.

Human oversight is non-negotiable but must be structured. Every topic in this report reaffirms that AI supports but does not replace human judgement in regulated fund administration activities. However, the report also makes clear that effective human oversight requires more than good intentions: it demands structured review processes (including the two-layer assurance framework for retrieval scope verification and substantive output quality), clear decision rights, documented escalation pathways, and role-specific training that equips staff to critically assess AI-generated content. Without these structures, human-in-the-loop becomes a nominal safeguard rather than an effective control.

Skills development and change management determine adoption success. The L3A AI Survey identified organizational resistance as a significant barrier to AI adoption. Topic 4 demonstrates that this resistance is best addressed through targeted, role-specific training grounded in real operational scenarios, transparent communication about the purpose and limitations of AI tools, and a culture that positions AI as an enabler of professional expertise rather than a threat to it. The AI literacy obligation under the EU AI Act (Article 4) provides the regulatory anchor for these efforts, but the practical implementation must go beyond compliance to foster genuine capability and confidence across the workforce.

Industry collaboration accelerates responsible adoption. The challenges identified across the four topics - regulatory complexity, data governance requirements, vendor management obligations, and the ManCo-PSF delegation boundary - are shared challenges. They are more effectively addressed through industry-level collaboration than through isolated, firm-by-firm responses. L3A is well positioned to continue this work: developing standardised contractual templates for ManCo-PSF AI governance, establishing shared best practices for AI vendor due diligence, and engaging with the CSSF on the proportionate application of the regulatory framework to AI deployment in fund administration.

This Consolidated Report represents a first, substantive step in translating the findings of the L3A AI Survey into actionable industry guidance. It does not claim to provide definitive answers to every question - the AI landscape evolves rapidly, and the regulatory framework continues to develop. What it does provide is a rigorous, practitioner-driven foundation on which Luxembourg fund administrators can build their AI strategies with confidence: anchored in the correct legal and regulatory baseline, designed for the right use-case architecture, and calibrated to the organization’s actual operational profile and ambitions.

The L3A AI Working Group invites all members to engage with this report, apply its recommendations to their own AI initiatives, and contribute to the ongoing development of industry standards that support safe, compliant, and effective AI adoption in Luxembourg fund administration.