GUARD GRC

Regulator Response Pack

DIFC Information Request

Prepared in response to a DIFC Commissioner of Data Protection / DFSA information request.

ORGANISATION
Meridian Capital Partners

SECTOR
financial-services

PRIMARY REGULATOR
DFSA

JURISDICTIONS
ae_difc, eu

PACK ID
2d9f8dbd-4317-4c4d-adab-7aae4474fbc0

GENERATED
2026-07-10T11:54:13.593Z by Sarah Al Mansoori

Scope: this pack is a point-in-time compilation of Meridian Capital Partners's AI-governance records as held in the GUARD GRC governance register at the generation timestamp above ("as at 2026-07-10T11:54:13.593Z"). The completeness manifest lists every requested item and its inclusion state — including items deliberately provided separately by the organisation. Nothing in this pack was newly generated for the pack: all embedded documents are reproduced verbatim from the register, each with a SHA-256 integrity hash (Integrity section).

Executive summary

AI-composed from the pack manifest at generation time; the underlying records are unmodified. Review before submission.

This executive summary accompanies a response pack prepared by Meridian Capital Partners for the Dubai Financial Services Authority under the DIFC Information Request template, addressing requirements applicable to its operations in the DIFC and EU jurisdictions. The pack reflects the state of the organisation's governance records as of 10 July 2026, 11:54:13 UTC.

The pack includes, in full, the following documents: a privacy notice (approved), an AI transparency statement (approved), and a data protection impact assessment (approved), each reproduced in Appendix B; an AI systems register comprising three active systems, one of which is classified; two completed AI risk assessments, reproduced in full in Appendix A; a record of the accountability and approval structure, listing three active members operating under a single-approval mode; and a governance audit trail covering the preceding twelve months, recording 38 consequential events.

Two items are included on a partial basis. AI governance policies and procedures comprise two approved policies embedded in full, alongside a third policy that remains in draft/review and is listed as pending but not embedded. Purpose and processing mapping is addressed through the DPIA and a per-regime classification register of seven rows, with a dedicated ROPA artefact identified as planned but not yet produced.

ADM attestations under DIFC Regulation 10.2.2 are not available within this pack; the manifest indicates these have not yet been generated and can be produced on demand via GUARD Attestations. Four further items are to be provided separately by the organisation from its own systems rather than through this pack: system design documentation, data-flow and architecture diagrams, DSAR request logs, and a model vendor security whitepaper concerning Anthropic.

All documents reproduced within the pack are drawn verbatim from the organisation's governance register and are accompanied by integrity hashes to support verification of their content against that register.

Contents

Completeness manifest

Checklist template: DIFC Information Request (difc_information_request). State as at 2026-07-10T11:54:13.593Z.

Privacy notice
Included
1 document (sign-off: approved) — full text in Appendix B.

AI transparency statement
Included
1 document (sign-off: approved) — full text in Appendix B.

AI systems register / inventory
Included
3 active systems (1 classified).

AI governance policies and procedures
Included — partial
2 approved policies embedded in full; 1 in draft/review (listed as pending, not embedded).

AI risk assessments
Included
2 completed assessments — full reports in Appendix A.

Data protection impact assessments (DPIA)
Included
1 document (sign-off: approved) — full text in Appendix B.

ADM attestations (DIFC Reg 10.2.2)
Not available
Not yet generated — available on demand via GUARD Attestations.

Purpose / processing mapping
Included — partial
Covered via DPIA (1) and the per-regime classification register (7 rows); a dedicated ROPA artefact is planned.

Accountability & approval structure
Included
3 active members; approval mode: single.

Governance audit trail (last 12 months)
Included
38 consequential events over 12 months.

System design documentation
Provided separately by the organisation
To be provided directly by the organisation from its own systems.

Data-flow / architecture diagrams
Provided separately by the organisation
To be provided directly by the organisation from its own systems.

DSAR request logs
Provided separately by the organisation
To be provided directly by the organisation from its own systems.

Model vendor security whitepaper (Anthropic)
Provided separately by the organisation
To be provided directly by the organisation (added at generation).

How this pack was assembled

• Every item was drawn live from the GUARD GRC governance register at the generation timestamp — nothing was created, regenerated or edited for this pack.

• The executive summary is the single exception: it is a narrative composition over the completeness manifest (AI-composed, flagged in place).

• Embedded documents are reproduced verbatim; each carries a SHA-256 hash of its stored content in the Integrity section.

• Sign-off states, approval chains and the audit-trail extract come from the append-only workflow ledger (guard.signoffs / guard.workflow_events), whose write paths are restricted to the platform's state machine.

• Per-subject audit packets previously issued from the register are referenced below by id and hash; their PDFs are independent tamper-evident records.

Previously issued audit packets

ef31ed4e · policy · 2026-07-08 17:38:14 UTC · SHA-256 c9bff9496e42b8eb…

AI systems register

Client Advisory Chatbot

Vendor: OpenAI (GPT-4o via API) · Type: genai_assistant · Lifecycle: production · Risk class: unclassified · User impact: — · Role: — · Not yet assessed

Client-facing conversational assistant on the Meridian digital advisory channel. Answers product and portfolio questions, provides market commentary, and routes clients to human advisers for regulated advice. Deployed on the retail web portal and mobile app; all conversations logged; escalation rules route any investment-advice intent to a licensed adviser.

CreditLens

Type: scoring_model · Lifecycle: production · Risk class: high · User impact: high · Role: provider · Last assessed: 2026-06-18 17:39:24 UTC

In-house credit and counterparty scoring model used in lending, margin and counterparty-limit decisions. Gradient-boosted model trained on internal repayment history, bureau data and transaction features. Scores feed the credit committee workflow; adverse scores can lead to declined facilities or reduced limits for retail and SME clients, including EU-domiciled clients of the Frankfurt branch.

Data categories: personal, financial

Sign-off: approved (approved 2026-06-20 17:39:24 UTC), expires 2027-06-20 17:39:24 UTC

• EU_AI_ACT: high — Assessment classifies CreditLens as high-risk under Annex III (credit scoring / access to essential financial services, Article 6(2)), affecting EU-domiciled clients via the Frankfurt branch. Conformity assessment, technical documentation (Art. 11/Annex IV), EU database registration (Art. 49) and human oversight (Art. 14) obligations apply but are only partially met.

• GDPR: gaps_identified — Personal and financial data of EU-domiciled clients is processed (CLS_04), with automated decisions producing legal/similarly significant effects, triggering Article 35 DPIA and Article 22 obligations. No DPIA has been completed (DPS_02_NO) and special-category safeguards and minimisation/retention are only partial, constituting identified gaps rather than full compliance.

• DIFC_REG10: gaps_identified — The system is registered in the DIFC autonomous systems register (Reg. 10.2 satisfied, DIFS_01_YES) and has a named accountable person (DIFS_03_YES), but fairness/non-discrimination testing required under Regulation 10.3 is not yet systematic (DIFS_02_PART), leaving gaps.

• UAE_PDPL: gaps_identified — UAE mainland operations trigger Federal PDPL Article 6 lawful basis and Article 21 impact assessment requirements; lawful basis is documented but no impact assessment equivalent to a DPIA has been completed and data minimisation/retention is only partially configured, per the report's UAE PDPL citations.

• CBUAE_AI_ML: tier_2 — CreditLens is a CBUAE-regulated lending activity (GOV_11_FS) with governance, explainability and outsourcing controls only partially met (GOV_12_PART) — indicating moderate but incomplete alignment with CBUAE AI/ML Guidance, consistent with a mid-tier risk categorization rather than fully compliant (tier_1) or seriously deficient (tier_3).

• NIST_AI_RMF: partial — The report notes NIST AI RMF as advisory with GOVERN/MAP/MEASURE/MANAGE functions applied but incomplete — e.g., value-vs-failure-cost tracking (Map 5) and incident response (Manage) are only partially implemented (GOV_09_PART, GOV_06_PART).

• ISO_42001: partial — Audit trail (Clause 8) is only partially reconstructable (GOV_02_PART) and decommissioning planning (Clause 8.4) exists only informally (GOV_10_PART), indicating partial alignment with ISO/IEC 42001 management system requirements.

TalentScreen CV Ranker

Vendor: TalentScreen FZ-LLC · Type: hr_screening · Lifecycle: pilot · Risk class: unclassified · User impact: — · Role: — · Not yet assessed

Third-party SaaS used by HR to rank inbound CVs against role profiles for shortlisting. Processes candidate personal data including employment history and education. Currently used for front-office and technology hires; shortlist decisions are reviewed by recruiters before rejection emails are sent.

Accountability & approval structure

Approval mode: single approver.

Team roster

Sarah Al Mansoori · owner · member since 2026-05-27
Omar Haddad · editor · member since 2026-05-28
Elena Rossi · viewer · member since 2026-05-28

Designated approvers

Omar Haddad (Chief Risk & Compliance Officer) — policy, action plan, assessment, attestation, ai system

Policies & procedures

AI Governance and Human Oversight Policy

Sign-off: approved 2026-06-05 17:39:24 UTC by Omar Haddad (Chief Risk & Compliance Officer); approval expires 2027-06-05 17:39:24 UTC
2026-06-04 17:39:24 UTC — Sarah Al Mansoori designated approver/reviewer
2026-06-05 05:39:24 UTC — Sarah Al Mansoori (owner): draft !' pending_approval
2026-06-05 17:39:24 UTC — Omar Haddad (approver): pending_approval !' approved
Implementation plan: completed 2026-06-28 17:39:24 UTC — completion certified

AI Governance and Human Oversight Policy

Organisation: Meridian Capital Partners
Effective Date: 8 July 2026
Version: 1.0 (Draft)
Classification: Internal
Policy Owner: Head of AI Governance (interim: Founder/CEO, pending formal appointment)

Prepared from organisation-provided information. The Policy Owner must verify the Scope section reflects the organisation's current state before approval.

1. Purpose

This policy governs the design, deployment, oversight, and risk classification of artificial intelligence systems used by Meridian Capital Partners, including AI systems the firm builds, systems it deploys from third parties, and general-purpose AI tools used informally by staff. It establishes the traceability, human oversight, training, and registration requirements necessary to meet Meridian's obligations under the EU AI Act (Articles 14 and 26), DIFC Data Protection Law and DIFC Regulation 10 (in particular Regulation 10.2 and 10.3.1(a)), and the NIST AI Risk Management Framework GOVERN function. This document constitutes Meridian's AI governance policy and is effective upon approval by the Policy Owner.

2. Scope

This policy applies to:

- CreditLens — Meridian's in-house credit and counterparty scoring model used in lending and margin decisions (EU AI Act Annex III §5 high-risk system).
- Advisory chatbot — the client-facing digital advisory channel providing retail investment guidance.
- Third-party CV-screening tool — used in HR recruitment (EU AI Act Annex III §4 high-risk system).
- Any general-purpose AI assistant or tool used by Meridian personnel in connection with company business, whether formally sanctioned or not.
- All Meridian personnel in the Dubai (DIFC) headquarters and Frankfurt branch, including employees, contractors, and personnel with AI oversight duties.
- All AI systems processing data of EU-domiciled clients, UAE-based clients, institutional clients, and employees, across both the DIFC and EU operating footprints.

This policy does not itself establish DFSA conduct-of-business rules or CBUAE model-risk methodology, which remain governed by their respective regulatory instruments; it operates alongside them.

3. Definitions

- AI System: Any software system, including CreditLens, the advisory chatbot, the CV-screening tool, and general-purpose AI assistants, that generates outputs such as predictions, recommendations, content, or decisions influencing Meridian's business or its clients/employees.
- High-Risk AI System: An AI system falling within EU AI Act Annex III, including credit scoring (§5) and employment/recruitment screening (§4).
- Human Oversight Personnel: Employees assigned responsibility for reviewing, validating, or intervening in AI system outputs before or after those outputs affect a client, employee, or business decision.
- Shadow AI: Use of general-purpose AI tools by staff without organisational approval, a defined permitted-tools list, or oversight, as currently practiced at Meridian.
- AI System Register: The standing inventory of Meridian's AI systems, their risk classification, owners, and status, maintained under Section 4.4.

4. Policy Statements

4.1 Human Oversight and Review Quality (CreditLens, Advisory Chatbot)

Human Oversight Personnel must apply a documented review standard before any CreditLens output is used in a

lending or margin decision, and before any advisory chatbot output is relied upon for retail investment guidance beyond pre-approved response templates. The review standard requires: (a) verification that the AI output is consistent with underlying data inputs; (b) documented sign-off recording the reviewer's identity and date; (c) escalation of any output that a reviewer cannot explain or validate to the Head of AI Governance before it is acted upon. The Head of AI Governance conducts quarterly quality-assurance spot-checks on a sample of reviewed decisions, examining whether reviewers applied the scrutiny criteria consistently. Human Oversight Personnel retain the standing authority to override, reverse, or disregard any CreditLens or chatbot output.

Regulatory basis: EU AI Act Art. 14 (human oversight design and effectiveness).
Responsible party: Head of AI Governance; day-to-day review performed by designated Human Oversight Personnel in Credit Risk and Client Advisory functions.

4.2 Training for AI Oversight Personnel

All personnel designated as Human Oversight Personnel for CreditLens, the advisory chatbot, or the CV-screening tool complete AI oversight training before assuming oversight duties and on an annual refresher basis thereafter. Training covers: the capabilities and limitations of the relevant AI system, the scrutiny criteria in Section 4.1, override and escalation procedures, and recognised failure modes (bias, hallucination, drift). No employee performs unsupervised human oversight of a high-risk AI system without having completed this training. The Head of AI Governance maintains training completion records for each Human Oversight Personnel member.

Regulatory basis: EU AI Act Art. 14 (competence, training and authority of persons assigned human oversight).
Responsible party: Head of AI Governance.

4.3 Decision Traceability and Logging

Every decision produced or materially informed by CreditLens, the advisory chatbot, or the CV-screening tool is logged at the point of use. The log captures, at minimum: the model or system version used, the data inputs (data lineage) relied upon, the identity of any human reviewer and their approval or override, and the rationale for the final decision where a human reviewer intervenes. Logs are retained in a format that allows the Head of AI Governance, the Commissioner (DIFC), or another competent authority to reconstruct how a given decision was reached. Personnel must not use an AI system in a manner that bypasses logging.

Regulatory basis: EU AI Act Art. 26 (deployer obligations, including record-keeping and traceability); DIFC Regulation 10.3.1(a) (Accountability — mechanisms to ensure responsibility for AI system outcomes).
Responsible party: Head of AI Governance, supported by the system owner of each logged AI system (Credit Risk lead for CreditLens; Client Advisory lead for the chatbot; HR lead for the CV-screening tool).

4.4 AI System Register and Risk Classification

Meridian maintains a standing AI System Register recording, for every AI system in use (including CreditLens, the advisory chatbot, and the CV-screening tool): the system name, archetype (builder/deployer), business function, data subjects affected, applicable risk classification under DIFC Regulation 10 and EU AI Act Annex III, and the accountable system owner. Any new AI system, and any general-purpose AI tool approved for business use under Section 4.5, is added to the Register before it is used in production. The Head of AI Governance classifies each system's risk tier at the point of registration and re-classifies it upon material change to its function or data processing.

Regulatory basis: DIFC Regulation 10 (risk classification and register obligations); NIST AI RMF GOVERN 1.6 (inventory of AI systems).
Responsible party: Head of AI Governance.

4.5 Governance of Informal ("Shadow") AI Use

Employees may use general-purpose AI assistants for Meridian business only where the tool appears on the Head of AI Governance's approved-tools list. Client data, employee personal data, and confidential firm information must not be entered into any AI tool that is not on the approved-tools list. Any employee currently using an unapproved AI tool for business purposes must cease that use or seek approval, in accordance with Annex A. Approved tools are added to the AI System Register under Section 4.4.

Regulatory basis: DIFC Regulation 10.2 (governance obligations for Personal Data processing by AI systems); NIST AI RMF GOVERN 1.6.
Responsible party: Head of AI Governance.

4.6 Governance Accountability

The Head of AI Governance is accountable for the operation of this policy, for the AI System Register, for decision-logging integrity, and for oversight training. Where AI-related risk exceeds the Head of AI Governance's authority to resolve (for example, a decision to suspend a high-risk system), the matter is escalated to the CEO for decision.

Regulatory basis: DIFC Regulation 10.2; NIST AI RMF GOVERN 2.1 (accountability structures and clear roles).
Responsible party: Head of AI Governance; CEO for escalations.

5. Roles and Responsibilities

Responsibility · Accountable Role · Note
Overall policy ownership and enforcement · Head of AI Governance · Currently held by Founder/CEO; delegate to a named Head of AI Governance once appointed (see Annex A)
Human oversight of CreditLens outputs · Credit Risk Lead (Human Oversight Personnel) · Reports quality issues to Head of AI Governance
Human oversight of advisory chatbot outputs · Client Advisory Lead (Human Oversight Personnel) · Reports quality issues to Head of AI Governance
Human oversight of CV-screening tool outputs · HR Lead (Human Oversight Personnel) · Reports quality issues to Head of AI Governance
Maintenance of AI System Register · Head of AI Governance · Updated on system change or addition
Decision-log integrity and retention · System owners (Credit Risk, Client Advisory, HR leads) · Overseen by Head of AI Governance
Approved AI tools list (shadow AI) · Head of AI Governance · Reviewed quarterly
Training delivery and records · Head of AI Governance · Annual refresher cycle
Escalation of unresolvable AI risk · CEO · Triggered by Head of AI Governance referral

6. Compliance Monitoring

The Head of AI Governance performs the following ongoing checks:

• Quarterly: QA spot-check sample of CreditLens and chatbot decisions against the review standard in Section 4.1; review of decision logs for completeness (model version, data lineage, reviewer, rationale); review of the approved AI tools list against actual employee usage.

• Semi-annually: Confirmation that all active Human Oversight Personnel hold current training records; review of the AI System Register for completeness and accuracy against actual systems in use.

• Annually: Full re-classification review of each registered AI system's risk tier; review of override and escalation logs to identify recurring issues requiring policy revision.

Evidence retained: QA spot-check results, training completion records, the AI System Register (current and historical versions), and decision logs, each retained for a minimum of six years or as required by applicable DIFC, EU, or DFSA record-keeping rules, whichever is longer.

7. Review and Update

This policy is reviewed at least annually by the Head of AI Governance. An early review is triggered by: introduction of a new AI system or material change to an existing one (including CreditLens model retraining or CV-screening tool provider change); a regulatory finding, audit observation, or incident involving an AI system; a material change to the EU AI Act, DIFC Regulation 10, or DFSA rules affecting Meridian's AI use; or appointment of a formal Head of AI Governance requiring role recalibration.

Annex A — Implementation Plan

Gap Addressed · Action · Owner · Deadline
No formally appointed AI governance role · Formally appoint a named Head of AI Governance (or confirm interim CEO holder in writing) · CEO · Immediate
AI System Register incomplete · Populate the AI System Register with full entries for CreditLens, advisory chatbot, and CV-screening tool, including risk classification · Head of AI Governance · 30 days
Decision-logging mechanism not yet built · Stand up the technical logging mechanism (model version, data lineage, reviewer, rationale) for CreditLens and the chatbot · Head of AI Governance + system owners · 90 days
Inconsistent human review quality · Document the scrutiny criteria referenced in Section 4.1 as a working checklist and distribute to Human Oversight Personnel · Head of AI Governance · 30 days
No formal oversight training · Design and deliver first training session to all current Human Oversight Personnel · Head of AI Governance · 90 days
Shadow AI use uncontrolled · Survey staff to identify current unapproved AI tool usage; issue approved-tools list and communicate Section 4.5 to all staff · Head of AI Governance · 30 days
CV-screening tool provider status unclear · Confirm with third-party vendor whether Meridian is deployer or co-developer under EU AI Act Annex III §4, and obtain vendor conformity documentation · Head of AI Governance · 90 days
CreditLens conformity status unclear · Commission a model validation/conformity assessment consistent with EU AI Act high-risk obligations and CBUAE model-risk expectations · Head of AI Governance · 90 days

Annex B — Items for the Policy Owner

• Confirm whether Meridian is the "deployer" or "provider/co-developer" of the third-party CV-screening tool under the EU AI Act, as this determines which Annex III obligations (provider vs. deployer) apply directly to Meridian.

• Obtain and review the full text of DFSA Rulebook provisions applicable to AI-assisted investment advisory and

discretionary mandate activity, which was not available in the material used to draft this policy.
• Verify whether CreditLens or the advisory chatbot processes any special-category or sensitive personal data under GDPR or DIFC Regulation 10, as this may trigger additional obligations not addressed here.
• Confirm the cross-border data transfer mechanism used between the DIFC and EU (Frankfurt) offices (e.g., adequacy decision, SCCs, or DIFC transfer agreement under DIFC Data Protection Law Articles 26–27), as this affects data flows underlying CreditLens and the chatbot but was not resolved in the source material.
• Verify existing CBUAE model-risk management documentation for CreditLens, if any, to confirm alignment with Section 4.6 escalation triggers.

AI Risk, Incident Response, and Data Breach Management Policy

Sign-off: approved 2026-06-10 17:39:24 UTC by Omar Haddad (Chief Risk & Compliance Officer); approval expires 2026-08-02 17:39:24+00:00

2026-06-09 17:39:24 UTC — Sarah Al Mansoori designated approver/reviewer
2026-06-10 05:39:24 UTC — Sarah Al Mansoori (owner): draft !' pending_approval
2026-06-10 17:39:24 UTC — Omar Haddad (approver): pending_approval !' approved

AI Risk, Incident Response, and Data Breach Management Policy

Organisation: Meridian Capital Partners
Effective Date: 8 July 2026
Version: 1.0 (Draft)
Classification: Internal
Policy Owner: AI Governance Lead (interim: Founder/CEO)

Prepared from organisation-provided information. The Policy Owner must verify the Scope section reflects the organisation's current state before approval.

1. Purpose

This policy governs how Meridian Capital Partners identifies, assesses, and manages financial exposure arising from its use of artificial intelligence, how it detects and responds to AI-related incidents, and how it responds to personal data breaches. It gives effect to the Company's obligations under the EU AI Act (Article 99, Article 26(5)), the EU GDPR (Article 33), the DIFC Data Protection Law (Article 42 and related breach provisions) and DIFC Regulation 10, and reflects the risk-management expectations of the CBUAE Guidance Note on Consumer Protection and Responsible Adoption of AI/ML. This policy consolidates what were previously separate, incomplete practices into a single standing framework covering financial risk exposure, AI incident response, and data breach response.

2. Scope

This policy applies to:

• All AI systems operated, built, or deployed by Meridian Capital Partners, specifically: CreditLens (in-house credit and counterparty scoring model used in lending and margin decisions), the Advisory chatbot (client-facing digital advisory channel), and the third-party CV-screening tool (HR recruitment).
• All uses of general-purpose AI assistants by employees, including informal or unapproved ("shadow AI") use, pending formal authorisation under Section 4.4.
• All personal data processed by the Company, including data of EU-domiciled clients (via the Frankfurt branch), UAE-based clients, institutional clients, and employees (including CV-screening candidates).
• All employees, contractors, and third-party service providers of Meridian Capital Partners operating from the Dubai (DIFC) headquarters or the Frankfurt branch.
• All jurisdictions in which the Company operates: the DIFC (as the Company's home regulatory base) and the European Union (via the Frankfurt branch and EU client base), together with any applicable UAE federal requirements.

3. Definitions

• AI Incident: Any event involving an AI system operated or deployed by the Company that causes, or has the potential to cause, harm to a data subject, client, employee, or the Company, including erroneous outputs, model failure, bias-driven decisions, unavailability, or security compromise.
• Personal Data Breach: A breach of security leading to the accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to, personal data, as defined under GDPR Article 4(12) and the DIFC Data Protection Law.
• High-Risk AI System: An AI system falling within Annex III of the EU AI Act, including CreditLens (credit scoring, Annex III §5) and the CV-screening tool (employment screening, Annex III §4).
• Financial Exposure: The Company's quantified potential liability arising from AI-related regulatory penalties, client compensation, remediation costs, and operational disruption.

4. Policy Statements

4.1 AI Financial Exposure and Risk Management Framework

The Company maintains a documented AI risk management framework covering every AI system in use, including CreditLens, the Advisory chatbot, and the CV-screening tool. For each system, the framework records: intended purpose, risk classification (including EU AI Act Annex III status where applicable), potential financial exposure (regulatory fines, client remediation, litigation, operational loss), and mitigating controls. The AI Governance Lead updates the financial exposure estimate whenever a new AI system is introduced, an existing system's use case materially changes, or following any AI incident.

Financial exposure estimates account for the maximum administrative fine exposure under applicable law, including EU AI Act Article 99 (fines up to 7% of worldwide annual turnover or EUR 35,000,000, whichever is higher, for the most serious infringements, scaled proportionately by infringement category) and GDPR Article 83 (fines up to 4% of worldwide annual turnover or EUR 20,000,000, whichever is higher). The AI risk management framework is integrated into, and does not duplicate, the Company's existing operational and conduct risk framework, consistent with CBUAE guidance that AI risk assessments should inform and be informed by overall risk appetite and controls.

Regulatory basis: EU AI Act Art. 99; GDPR Art. 83; CBUAE Guidance Note (risk appetite integration). Responsible party: AI Governance Lead, with escalation to the CEO for exposure estimates exceeding a materiality threshold set by the CEO.

4.2 AI Incident Response

The Company maintains a standing AI incident response procedure applicable to all AI systems in scope of this policy. The procedure requires:

- Detection: Each AI system owner (CreditLens: Credit Risk function; Advisory chatbot: Digital Channels function; CV-screening tool: HR function) monitors their system for erroneous, biased, or harmful outputs and reports suspected incidents to the AI Governance Lead within 24 hours of detection.
- Triage and containment: The AI Governance Lead classifies each reported incident by severity and, where warranted, directs the system owner to suspend or restrict the AI system's operation pending investigation.
- Regulatory disclosure: Where an AI incident constitutes a "serious incident" affecting an EU high-risk AI system, or otherwise triggers a regulatory notification obligation, the AI Governance Lead notifies the relevant authority without undue delay, consistent with EU AI Act Article 26(5) obligations on deployers of high-risk AI systems.
- Record-keeping: The AI Governance Lead documents every AI incident in an AI incident log, recording the nature of the incident, affected systems and individuals, containment measures, root cause, and corrective action taken.
- Post-incident review: Following resolution of any incident classified as Moderate or above, the AI Governance Lead conducts a documented post-incident review and updates the relevant system's risk assessment and controls accordingly.

Regulatory basis: EU AI Act Art. 26(5); DIFC Data Protection Law (breach documentation obligations). Responsible party: AI Governance Lead (incident coordination); individual System Owners (detection and first-line escalation).

4.3 Personal Data Breach Response

The Company maintains a documented personal data breach response procedure covering detection, assessment, notification, and remediation. Under this procedure:

- Any employee who becomes aware of a suspected personal data breach notifies the Data Protection Lead immediately and in any event within 24 hours.
- The Data Protection Lead assesses whether the breach is likely to result in a risk to the rights and freedoms of affected data subjects and determines applicable notification obligations by jurisdiction.
- Where the breach involves personal data of EU-domiciled data subjects and is likely to result in a risk to rights and freedoms, the Data Protection Lead notifies the competent EU supervisory authority without undue delay and, where feasible, within 72 hours of becoming aware of the breach, in accordance with GDPR Article 33. Where notification is not made within 72 hours, the notification records the reasons for the delay.
- Where the breach involves personal data subject to the DIFC Data Protection Law, the Data Protection Lead reports the breach to the DIFC Commissioner without undue delay, consistent with DIFC Regulation 10 breach-reporting requirements.
- Every notification includes, at minimum: the nature of the breach and approximate number of data subjects and records affected; the name and contact details of the Data Protection Lead; the likely consequences of the breach; and the measures taken or proposed to address the breach and mitigate its adverse effects. Where full information is not available at the time of notification, the Data Protection Lead provides it in phases without further undue delay.

• The Company notifies affected data subjects directly where the breach is likely to result in a high risk to their rights and freedoms, using the notification content standard set out above.

• The Data Protection Lead documents every personal data breach in writing, recording the facts, effects, and remedial action taken, in a form sufficient to allow the relevant supervisory authority to verify compliance.

• The Company maintains a standard notification template and a maintained contact list of relevant supervisory authorities (DIFC Commissioner and the competent EU supervisory authority for the Frankfurt branch) to enable timely notification.

Regulatory basis: GDPR Art. 33; DIFC Data Protection Law (breach reporting and documentation provisions); DIFC Regulation 10.
Responsible party: Data Protection Lead.

4.4 Shadow AI Use

Pending the outcome of the acceptable-use assessment in Annex A, employees using general-purpose AI assistants for business purposes must not input client personal data, employee personal data, or confidential firm information into any AI tool that has not been reviewed and approved by the AI Governance Lead. This restriction applies with immediate effect.

Regulatory basis: GDPR Art. 33 and general data minimisation principles underpinning DIFC Data Protection Law; EU AI Act risk-management principles.
Responsible party: AI Governance Lead; all employees.

5. Roles and Responsibilities

Responsibility · Accountable Role · Notes
Overall ownership of this policy · Policy Owner (AI Governance Lead) · Currently held by Founder/CEO pending formal appointment (see Annex A)
AI financial exposure estimation and updates · AI Governance Lead · Escalates material exposure to CEO
AI incident detection (per system) · System Owners: Credit Risk (CreditLens), Digital Channels (Advisory chatbot), HR (CV-screening tool) · First-line reporting duty
AI incident triage, coordination, regulatory disclosure · AI Governance Lead ·
Personal data breach assessment and regulatory notification · Data Protection Lead · Currently held by Founder/CEO pending formal appointment (see Annex A)
Data subject notification · Data Protection Lead ·
Maintenance of AI incident log and breach documentation · AI Governance Lead / Data Protection Lead (respective logs) ·
Approval of GenAI tools for staff use · AI Governance Lead ·

6. Compliance Monitoring

The Policy Owner conducts the following checks:

• Quarterly: Review of the AI financial exposure estimate for each in-scope AI system, confirming it reflects current use, volume, and any regulatory developments. Evidence retained: updated exposure register with revision history.

• Quarterly: Review of the AI incident log to confirm all reported incidents were triaged, documented, and (where applicable) escalated within required timeframes. Evidence retained: incident log with resolution status.

• Semi-annually: Test of the personal data breach notification procedure (tabletop exercise) confirming the 72-hour GDPR notification pathway and DIFC Commissioner reporting pathway are current and functional. Evidence retained: exercise report and any identified corrective actions.

• Annually: Confirmation that the AI incident log and breach register are complete, that notification templates remain accurate (contact details, authority names), and that CreditLens and the CV-screening tool retain current risk classifications. Evidence retained: annual compliance attestation signed by the Policy Owner.

• Ongoing: Monitoring of shadow AI usage reports and enforcement of the interim restriction in Section 4.4.

7. Review and Update

This policy is reviewed at least annually by the Policy Owner. An early review is triggered by: any AI incident classified Moderate or above; any personal data breach requiring regulatory notification; introduction of a new AI system or material change to an existing system's use case; a change in applicable law (including new EU AI Act implementing acts or DIFC Regulation amendments); or a material change to the Company's operating footprint (e.g., new branch or jurisdiction).

Annex A — Implementation Plan

Gap Addressed · Action · Owner · Deadline
AI financial exposure estimate incomplete · Populate the AI risk/exposure register for CreditLens, Advisory chatbot, and CV-screening tool with current use, volume, and quantified exposure estimates · AI Governance Lead (interim:

CEO) · 30 days
AI risk management framework incomplete · Complete first full risk assessment for CreditLens (credit scoring, Annex III §5) and CV-screening tool (Annex III §4), including EU AI Act conformity gap analysis · AI Governance Lead (interim: CEO) · 90 days
No formal AI incident response plan · Stand up the AI incident log; brief System Owners (Credit Risk, Digital Channels, HR) on detection and escalation duties under Section 4.2 · AI Governance Lead (interim: CEO) · 30 days
No formal AI incident response plan · Conduct first tabletop exercise of the AI incident response procedure · AI Governance Lead (interim: CEO) · 90 days
Partial data breach response plan · Finalize and distribute breach notification template; confirm current DIFC Commissioner and EU supervisory authority contact details · Data Protection Lead (interim: CEO) · 30 days
Partial data breach response plan · Conduct first tabletop exercise of the 72-hour GDPR / DIFC breach notification pathway · Data Protection Lead (interim: CEO) · 90 days
Governance roles not formally appointed · Formally appoint an AI Governance Lead and Data Protection Lead (may be combined initially), distinct from the CEO where feasible · CEO · 90 days
Shadow AI use unmanaged · Issue staff communication on interim GenAI restriction (Section 4.4); begin review of commonly used GenAI tools for formal approval · AI Governance Lead (interim: CEO) · Immediate
Cross-border transfer mechanism unclear · Confirm and document the transfer mechanism used between the DIFC and Frankfurt branch (adequacy decision, SCCs, or DIFC-approved clauses) · Data Protection Lead (interim: CEO) · 90 days

Annex B — Items for the Policy Owner

• Confirm whether Meridian Capital Partners is the "provider" or "deployer" of the third-party CV-screening tool under the EU AI Act, as this determines which Article 26/Article 16 obligations apply directly (current information does not establish provider identity).

• Confirm whether CreditLens or the Advisory chatbot processes special-category data under GDPR Article 9; this policy does not currently address special-category data handling and may require a supplementary annex if confirmed.

• Obtain and review DIFC Regulation 10, Section 10.2 and 10.3 in full detail before final approval, to confirm all deployer/operator obligations for autonomous systems (human-intervention algorithms, registers, risk and impact assessments) are reflected in Section 4.

• Verify whether the DFSA Rulebook imposes additional incident-notification or model-risk obligations on the discretionary mandate and retail advisory activities, given the confirmed DFSA-regulated status, and whether these should be incorporated as a further policy section.

• Verify current AUM/revenue bands to confirm whether Meridian falls within or outside the 50-person threshold exemption referenced in DIFC Data Protection Regulations Regulation 2.2 (the Company's ~180 employees suggests the exemption does not apply, but this should be confirmed against "High Risk Processing Activities" classification).

• Confirm CBUAE Model Management Standards (2022) documentation status for CreditLens, referenced in the CBUAE Guidance Note, as this may require a dedicated model-risk annex beyond the scope of this policy.

Pending (not embedded)

• AI Governance and Risk Management Policy — pending approval

Attestations & impact assessments

Current (latest, non-superseded) governance artefacts. Full documents are reproduced in Appendix B.

Privacy Notice — Meridian Capital Partners

Privacy Notice · Company-wide · generated 2026-07-01 · next review 2027-07-01
Sign-off: approved (approved 2026-07-03 19:27:38 UTC) · content SHA-256 1bafc4a2de9762cf…

AI Transparency Statement — Meridian Capital Partners

AI Transparency Statement · Company-wide · generated 2026-07-01 · next review 2027-07-01
Sign-off: approved (approved 2026-07-03 19:27:38 UTC) · content SHA-256 192b3821e72218f6…

Data Protection Impact Assessment: CreditLens Credit and Counterparty Scoring Model

Data Protection Impact Assessment · CreditLens · generated 2026-06-24
Sign-off: approved (approved 2026-06-25 17:39:24 UTC) · content SHA-256 d3f5c3afd93028fb…

Governance audit trail (extract)

38 consequential events over 12 months.

2026-07-07 17:39:24 UTC · Sarah Al Mansoori · policy · sign-off: draft !’ pending_approval — “Retention section aligned with DIFC/ DFSA as requested.”
2026-07-05 17:39:24 UTC · Omar Haddad · action plan · sign-off: pending_approval !’ approved
2026-07-03 19:27:38 UTC · Omar Haddad · attestation · sign-off: pending_approval !’ approved
2026-07-03 17:39:24 UTC · Sarah Al Mansoori · policy · edit unapproved
2026-07-03 14:39:38 UTC · Sarah Al Mansoori · attestation · sign-off: draft !’ pending_approval
2026-07-03 09:51:38 UTC · Sarah Al Mansoori · attestation · sign-off designated (Chief Risk & Compliance Officer)
2026-07-03 05:03:38 UTC · Omar Haddad · attestation · sign-off: pending_approval !’ approved
2026-07-03 00:15:38 UTC · Sarah Al Mansoori · attestation · sign-off: draft !’ pending_approval
2026-07-02 19:27:38 UTC · Sarah Al Mansoori · attestation · sign-off designated (Chief Risk & Compliance Officer)
2026-06-30 08:03:24 UTC · Sarah Al Mansoori · action plan · sign-off: draft !’ pending_approval
2026-06-29 17:39:24 UTC · Omar Haddad · policy · sign-off: pending_approval !’ approved
2026-06-29 17:39:24 UTC · Omar Haddad · action plan · sign-off: pending_approval !’ approved — “Execution verified — all controls in place.”
2026-06-26 11:39:24 UTC · Sarah Al Mansoori · action plan · plan completed
2026-06-25 17:39:24 UTC · Omar Haddad · attestation · sign-off: pending_approval !’ approved
2026-06-25 17:39:24 UTC · Sarah Al Mansoori · policy · sign-off: draft !’ pending_approval
2026-06-25 05:39:24 UTC · Sarah Al Mansoori · attestation · sign-off: draft !’ pending_approval
2026-06-24 22:27:24 UTC · Sarah Al Mansoori · action plan · sign-off designated (Chief Risk & Compliance Officer)
2026-06-24 17:39:24 UTC · Sarah Al Mansoori · attestation · sign-off designated (Chief Risk & Compliance Officer)
2026-06-23 05:39:24 UTC · Sarah Al Mansoori · action plan · plan completion submitted
2026-06-21 17:39:24 UTC · Sarah Al Mansoori · policy · sign-off designated (Chief Risk & Compliance Officer)
2026-06-20 17:39:24 UTC · Omar Haddad · ai system · sign-off: pending_approval !’ approved
2026-06-20 05:39:24 UTC · Sarah Al Mansoori · ai system · sign-off: draft !’ pending_approval
2026-06-19 23:39:24 UTC · Omar Haddad · action plan · sign-off: pending_approval !’ approved
2026-06-19 17:39:24 UTC · Sarah Al Mansoori · ai system · sign-off designated (Chief Risk & Compliance Officer)
2026-06-19 17:39:24 UTC · Omar Haddad · assessment · sign-off: pending_approval !’ approved
2026-06-19 12:51:24 UTC · Omar Haddad · policy · sign-off: pending_approval !’ approved
2026-06-19 05:39:24 UTC · Sarah Al Mansoori · assessment · sign-off: draft !’ pending_approval
2026-06-18 17:39:24 UTC · Sarah Al Mansoori · assessment · sign-off designated (Chief Risk & Compliance Officer)
2026-06-16 17:39:24 UTC · Sarah Al Mansoori · action plan · sign-off: draft !’ pending_approval
2026-06-14 03:15:24 UTC · Sarah Al Mansoori · policy · sign-off: draft !’ pending_approval
2026-06-13 11:39:24 UTC · Sarah Al Mansoori · action plan · sign-off designated (Chief Risk & Compliance Officer)
2026-06-10 05:39:24 UTC · Omar Haddad · policy · sign-off: pending_approval !’ approved
2026-06-08 17:39:24 UTC · Sarah Al Mansoori · policy · sign-off designated (Chief Risk & Compliance Officer)
2026-06-06 23:39:24 UTC · Sarah Al Mansoori · policy · sign-off: draft !’ pending_approval
2026-06-03 17:39:24 UTC · Sarah Al Mansoori · policy · sign-off designated (Chief Risk & Compliance Officer)
2026-06-02 17:39:24 UTC · Omar Haddad · assessment · sign-off: pending_approval !’ approved
2026-06-02 05:39:24 UTC · Sarah Al Mansoori · assessment · sign-off: draft !’ pending_approval
2026-06-01 17:39:24 UTC · Sarah Al Mansoori · assessment · sign-off designated (Chief Risk & Compliance Officer)

Appendix A — Assessment reports (full text)

Meridian Capital Partners

CreditLens · completed 2026-06-18 17:39:24 UTC · content SHA-256 6a17f7a46c1a04bc…

AI System Assessment Report

System: Multi-system AI portfolio comprising (1) CreditLens — an in-house gradient-boosted credit and counterparty scoring model used in lending, margin, and counterparty-limit decisions; (2) a client-facing digital advisory chatbot for retail investment guidance; and (3) a third-party CV-screening tool used in HR recruitment. This assessment is scoped primarily to the credit scoring system based on the classification answers provided (Annex III credit scoring trigger, human-on-the-loop autonomy, legal-effect outputs), with contextual notes on the wider AI estate where the profile summary materially affects interpretation.
Client: Meridian Capital Partners
Assessment Date: 8 July 2026
Assessor: GUARD GRC
Jurisdictions: European Union/EEA (Frankfurt branch clients), DIFC (Dubai headquarters), UAE (mainland)
Risk Classification: High-Risk under EU AI Act (Annex III — credit scoring / access to essential financial services, Article 6(2))

Executive Summary

Meridian Capital Partners has completed a GUARD Navigator assessment for an AI system used in credit and counterparty scoring, which the organization has classified as high-risk under EU AI Act Annex III on the basis of its use in "credit scoring, insurance pricing, or access to essential services" (CLS_06). The system operates in human-on-the-loop mode, processes personal and financial data, and produces outputs with legal or similarly significant effects on customers — findings consistent with the profile summary describing CreditLens, an in-house gradient-boosted scoring model that feeds credit-committee decisions and can result in declined facilities or reduced limits for retail and SME clients, including EU-domiciled clients served from the Frankfurt branch. The company operates as Provider (it developed the system) across the EU, DIFC, and UAE mainland, triggering overlapping obligations under the EU AI Act, GDPR, DIFC Data Protection Law and Regulation 10, and UAE PDPL.

DATA QUALITY ADVISORY: The Navigator responses do not directly identify the system by name, but the supplementary profile summary and system record identify it as CreditLens. We have proceeded on the basis that the Navigator responses (CLS_01–CLS_07, EUS_, DPS_, GOV_, DIFS_) describe CreditLens specifically, consistent with its credit-scoring function, financial data processing, and legal-effect outputs. We note that the profile summary describes a materially broader AI estate — an advisory chatbot and a third-party CV-screening tool — neither of which appears to have been separately assessed through this Navigator session. The CV-screening tool independently triggers EU AI Act Annex III §4 (employment/recruitment screening) as a high-risk use case, and the advisory chatbot triggers transparency obligations under Article 50 and potentially DFSA conduct rules. These two systems are outside the scope of the responses received and are not scored in this report. We recommend running separate Navigator assessments for the advisory chatbot and the CV-screening tool, and flag the disclosed "shadow-AI" use of general-purpose AI assistants by employees (profile summary, Section 4) as a governance gap requiring immediate attention outside this scored assessment.

Overall compliance posture across the five GUARD categories is mixed: Governance (G) and Unauthorized Data Access (U) are rated YELLOW, Attrition of Skills (A) is rated YELLOW after contextual adjustment, and Reputation (R) and Dollar Drain (D) are rated GREEN. This report identifies 9 findings requiring remediation (0 RED, 9 YELLOW), against a backdrop of genuine compliance strengths — documented lawful basis, mapped cross-border transfer safeguards, working human oversight, DIFC system registration, and production monitoring — which are acknowledged in the category narratives below but are not counted as findings.

Regulatory Applicability

Regulation · Applies · Trigger · Key Obligations
EU AI Act (Reg. 2024/1689) · Yes · EU-domiciled clients affected by outputs (CLS_01); credit scoring selected as Annex III high-risk use (CLS_06); Provider role (CLS_02) · Articles 9-15 (risk management, data governance, technical documentation, accuracy/robustness), Article 16 (provider obligations), Article 43 (conformity assessment), Article 49 (EU database registration), Article 27 (FRIA where applicable to deployers)
GDPR (Reg. 2016/679) · Yes · Personal data of EU-domiciled clients (CLS_04, CLS_01_EU); Frankfurt branch serves EU clients · Article 6 (lawful basis), Article 9 (special categories), Article 35 (DPIA), Chapter V Articles 44-49 (cross-border transfers), Articles 15-22 (data subject rights, automated decision-making)
DIFC Data Protection Law (Law No. 5 of 2020) · Yes · DIFC-licensed entity (CLS_01_DIFC); DIFC headquarters · Article 26 (transfers out of DIFC — adequacy or Article 27 safeguards), general lawful processing obligations equivalent to controller/processor duties

DIFC Regulation 10 · Yes · Autonomous/semi-autonomous system operated by a DIFC entity · Regulation 10.2 (system register — satisfied), Regulation 10.3 (fairness, ethical design, non-discrimination), Regulation 10.4 (named accountable person — satisfied)
UAE Federal PDPL (Decree-Law 45/2021) · Yes · UAE mainland operations (CLS_01_UAE) · Article 6 (lawful basis), cross-border transfer conditions, sensitive data controls
CBUAE AI/ML Guidance · Yes · Regulated financial-services activity confirmed (GOV_11_FS — lending under CBUAE supervision) · Governance and accountability, model risk management, explainability, third-party/outsourcing controls, ongoing monitoring
NIST AI RMF · Advisory · In-house model development (CreditLens) and high-risk deployment · GOVERN, MAP, MEASURE, MANAGE functions applied to model risk and data governance
ISO/IEC 42001 · Advisory · Multi-jurisdictional AI deployment posture · AI management system — accountability, operational records, decommissioning planning

Note on non-personal data: CLS_04 confirms personal and financial data are processed by this system — Rule 2's "no personal data" carve-out does not apply here. All cross-border transfer and data protection findings below are assessed on the basis that personal data is genuinely in scope.

GUARD Risk Assessment

G — Governance: YELLOW

Finding: Meridian has a named, empowered system owner (GOV_01_YES) and the system is properly recorded in the DIFC autonomous systems register as required by Regulation 10.2 (DIFS_01_YES) — both genuine strengths. However, technical documentation is incomplete (EUS_02_PART), the audit trail for decisions and changes is only partially reconstructable (GOV_02_PART), EU database registration has not been completed despite being required (EUS_04_NO), and CBUAE model-risk expectations for explainability and outsourcing controls are only partially met (GOV_12_PART) despite this being a CBUAE-regulated lending activity (GOV_11_FS).

Regulatory basis: EU AI Act Article 11 and Annex IV (technical documentation), Article 49 (EU database registration for high-risk systems), Article 16 (provider obligations); DIFC Regulation 10.2 (system register); CBUAE AI/ML Guidance Note Section 2 (governance and accountability) and model risk management/outsourcing expectations.

What's missing:
1. Complete technical documentation meeting EU AI Act Annex IV requirements (system description, design specifications, data governance summary, known limitations) — currently partial and outdated (EUS_02_PART).
2. A reliable audit trail covering approvals, training data provenance, and change history — currently only partially reconstructable (GOV_02_PART).
3. EU database registration under Article 49 has not been completed (EUS_04_NO) — this is a binding pre-deployment obligation for high-risk systems.
4. CBUAE explainability documentation and outsourcing/third-party controls for model components are only partially in place (GOV_12_PART).

Remediation:
1. Complete and update the Annex IV technical documentation file for CreditLens — commission a structured documentation exercise capturing model architecture, training data lineage, intended purpose, performance metrics, and known limitations.
2. Implement a version-controlled change log capturing every model retraining, parameter change, and committee approval, linked to the named system owner.
3. Complete EU database registration for CreditLens as a high-risk AI system before further processing of EU-domiciled client data.
4. Produce a CBUAE-ready explainability pack (e.g., SHAP-based feature attribution outputs or equivalent) and document outsourcing/vendor due diligence for any third-party components feeding the model.

Effort: Medium
Timeline: 30-60 days

U — Unauthorized Data Access: YELLOW

Finding: Meridian has a documented lawful basis for processing (DPS_01_YES) and has mapped and safeguarded cross-border transfers of this system's data (DPS_04_YES) — both strong compliance indicators reflecting genuine control maturity given the DIFC-to-EU-to-UAE data flows implied by the Frankfurt branch structure. However, no DPIA has been completed despite one likely being required given the high-risk, legal-effect nature of the system (DPS_02_NO), special-category data safeguards are only partially implemented (DPS_03_PART), data minimisation and retention limits are only partially configured (DPS_05_PART), and data governance over training data provenance and representativeness is only partially understood (EUS_05_PART).

Regulatory basis: GDPR Article 35 (DPIA — mandatory for automated decision-making producing legal effects per Article 35(3)(a)); UAE PDPL Article 21 (impact assessment); GDPR Article 9 (special categories); GDPR Article 5 (data minimisation, storage limitation); UAE PDPL Article 5 (processing principles); EU AI Act Article 10 (data

governance for high-risk systems).

What's missing:
1. No DPIA has been conducted for a system that scores individuals and produces legal/significant effects (declined facilities, reduced limits) — this is very likely a mandatory trigger under GDPR Article 35(3)(a) and UAE PDPL Article 21.
2. Special-category data safeguards (explicit consent or documented exemption plus tighter controls) are only partially in place.
3. Data minimisation and retention limits on inputs, prompts, and logs are only partially configured.
4. Training/input data governance — relevance, representativeness, and provenance — is only partially documented, a gap directly relevant to bias risk in a credit-scoring context.

Remediation:
1. Conduct a full DPIA for CreditLens covering necessity and proportionality, risk to data subjects (including discriminatory outcome risk), and mitigating measures, before further processing of EU or UAE client data in lending decisions.
2. Complete a special-category data mapping exercise to confirm whether any sensitive data categories (e.g., health data referenced in insurance/lending contexts, or proxies correlating with protected characteristics) enter the model, and implement the stronger safeguards required where applicable.
3. Configure and document retention schedules and minimisation rules for all data fields, prompts, and logs used by the scoring model.
4. Complete a data governance record for training and input data — source, collection method, representativeness testing, and update frequency — to close the EU AI Act Article 10 gap.

Effort: Medium-High
Timeline: 30-60 days (DPIA); 60-90 days (data governance documentation)

A — Attrition of Skills: YELLOW

Navigator score: yellow (1 gap indicator: GOV_04_NO). Adjusted assessment: YELLOW — confirmed, not downgraded, because the identified gap concerns the effectiveness of oversight rather than its existence, and this is a material finding given the credit-decision context.

Finding: Human oversight is designed and formally operating for this system (EUS_03_YES), and a named accountable person exists at the DIFC entity level (DIFS_03_YES) — genuine, positive structural controls. However, Meridian candidly reported that human review of this system's outputs has become largely a rubber stamp in practice (GOV_04_NO), and staff training is limited to general AI awareness rather than system-specific competence in CreditLens's limitations and failure modes (GOV_03_PART). This combination — formal oversight structure paired with degraded practical effectiveness — is precisely the automation-bias risk that Article 14 is designed to prevent, and is a more serious finding than a simple missing-control gap because the appearance of compliance masks an operational failure.

Regulatory basis: EU AI Act Article 14 (human oversight must enable individuals to fully understand, monitor, and be able to intervene — not merely be present); Article 26 (deployer obligation to ensure oversight is exercised by competent, trained, and authorized persons); Article 4 (AI literacy).

What's missing:
1. Human reviewers of CreditLens outputs are not meaningfully exercising independent judgment — oversight has degraded into automation bias (GOV_04_NO).
2. Training for staff operating or relying on the system covers general AI awareness only, not CreditLens-specific limitations, failure modes, or escalation triggers (GOV_03_PART).

Remediation:
1. Redesign the credit committee review workflow to require documented, substantive justification for concurrence with (not just override of) scoring outputs on a sample basis, and introduce periodic "challenge sessions" where reviewers must justify decisions independent of the score.
2. Deliver CreditLens-specific training to all credit committee members and reviewers covering the model's known limitations, historical failure modes, edge cases, and defined escalation triggers, with completion tracked and refreshed periodically.

Effort: Medium
Timeline: 30-60 days

R — Reputation: GREEN

Finding: Affected customers are informed that AI is involved and have access to a working explanation channel (GOV_07_YES) — a strong transparency control directly addressing EU AI Act Article 50 and GDPR Articles 13-15/22 obligations. Supporting controls are progressing but incomplete: bias testing has occurred once without a recurring schedule (GOV_05_PART), incident response planning exists partially (a kill-switch capability exists, but the surrounding response process is improvised) (GOV_06_PART), robustness/accuracy testing evidence is thin

(EUS_06_PART), disclosure of AI involvement is inconsistent across touchpoints (EUS_07_PART), high-risk conformity work for the EU AI Act is underway but incomplete (EUS_01_PART), and DIFC fairness/non-discrimination evidence is not yet systematic (DIFS_02_PART).

Regulatory basis: EU AI Act Article 50 (transparency), Article 10 (bias examination), Article 15 (accuracy, robustness, cybersecurity), Articles 16/26/43 (conformity for high-risk systems); DIFC Regulation 10.3 (fairness and non-discrimination); GDPR Articles 13-15 and 22 (information/explanation rights); NIST AI RMF Manage function (incident response); NYC Local Law 144 §20-871 (bias audit, if applicable to any US-facing hiring use — noted for the separately-scoped CV tool, not CreditLens).

The Navigator scored this category GREEN based on zero net gap indicators after suffix counting, since none of the individual answers crossed into outright "NO" territory except where noted. We concur with GREEN as the overall rating given the strong transparency foundation (GOV_07_YES) and the fact that all remaining items are partial/in-progress rather than absent. However, the six PARTIAL items below represent real, trackable gaps that should not be treated as closed.

What's missing:
1. Bias testing was performed once with no recurring schedule (GOV_05_PART) — a single point-in-time test does not demonstrate ongoing fairness, particularly important given the credit-scoring use case's discrimination risk.
2. Incident response beyond the ability to stop the system is improvised, with no documented rollback procedure, incident owner, or affected-party contact channel (GOV_06_PART).
3. Accuracy, robustness, and cybersecurity testing evidence is thin relative to what a regulator would expect to see (EUS_06_PART).
4. AI disclosure is inconsistent across customer touchpoints (EUS_07_PART).
5. High-risk conformity assessment work under Articles 16/26/43 is underway but not complete (EUS_01_PART).
6. DIFC fairness/non-discrimination evidence is ad hoc rather than systematic (DIFS_02_PART).

Remediation:
1. Establish a recurring bias-testing schedule (e.g., quarterly) for CreditLens covering protected/proxy characteristics relevant to lending decisions, with results reported to the credit committee and retained as evidence.
2. Document a formal incident response plan: named incident owner, rollback/kill-switch procedure, escalation path, and a defined channel for affected customers to raise concerns about a credit decision.
3. Commission structured accuracy, robustness, and cybersecurity testing (including adversarial/edge-case testing) and retain the resulting evidence in a form presentable to a regulator.
4. Standardize AI-disclosure language and placement across every customer touchpoint where CreditLens outputs affect a decision.
5. Complete the in-progress EU AI Act conformity assessment work and set a firm completion date tracked by the system owner.
6. Systematize DIFC fairness/non-discrimination testing into the same recurring schedule as Item 1, with results retained for the DIFC Commissioner.

Effort: Medium
Timeline: 30-90 days

D — Dollar Drain: GREEN

Finding: This is a genuine strength area. GPAI-related obligations are covered where applicable (EUS_08_YES), and the system is monitored in production with defined metrics and alerts (GOV_08_YES), giving Meridian visibility into performance degradation before it becomes a regulatory or customer-facing incident. The remaining item is value-tracking: Meridian tracks only one of the two elements of value delivery versus failure-cost estimation (GOV_09_PART), and a decommissioning plan exists only informally, not in writing (GOV_10_PART).

Regulatory basis: NIST AI RMF Map 5 (impact and value assessment); ISO/IEC 42001 Clause 8.4 (decommissioning and cessation); EU AI Act Article 72 (post-market monitoring — satisfied).

What's missing:
1. Value-versus-failure-cost tracking is incomplete — Meridian tracks only one of the two dimensions (whether the system delivers intended value, or the cost exposure if it failed/were switched off) (GOV_09_PART).
2. A decommissioning plan for CreditLens has been discussed informally but is not documented (GOV_10_PART).

Remediation:
1. Complete the value-tracking framework by adding the missing dimension — document an estimate of financial and operational exposure if CreditLens were disabled or failed, alongside existing value-delivery tracking.
2. Draft a written decommissioning plan covering data disposition, dependent-process transition, and stakeholder notification requirements, and store it alongside the system's governance record.

Effort: Low

Timeline: 0-30 days

Clarification Required

· Question · User Response · Issue · Impact on Assessment
1 · System identity/scope · Navigator responses do not name the system; supplementary materials identify it as CreditLens, and describe two additional AI systems (advisory chatbot, CV-screening tool) not covered by these responses · The scope of this Navigator session is not explicitly confirmed against the broader AI estate described in the profile summary · This assessment is limited to the system described by the Navigator responses (interpreted as CreditLens). The advisory chatbot and CV-screening tool require separate assessment and are not scored here.
2 · DPS_03 — special category data safeguards · "Partially" · It is unclear which special-category data types (if any) are processed and which safeguards are missing · Treated conservatively as a partial gap; recommend a data mapping exercise to confirm scope (see Rule 2 general principle)
3 · Shadow-AI use (profile summary only, not a scored Navigator question) · Employees use general-purpose AI assistants informally, with no permitted-tools policy · This is a material governance gap disclosed outside the scored questionnaire · Not included in the GUARD category scores or gap summary below, since no corresponding Navigator question was answered. Flagged here for organizational attention.

Compliance Gap Summary

· Finding · Regulation · GUARD Category · Severity · Remediation Type · Remediation · Timeline
1 · No DPIA completed despite likely mandatory trigger · GDPR Art. 35; UAE PDPL Art. 21 · U · YELLOW · operational · Conduct a full DPIA covering necessity, proportionality, and discriminatory-outcome risk · 30-60 days
2 · EU database registration not completed for high-risk system · EU AI Act Art. 49 · G · YELLOW · operational · Complete EU database registration for CreditLens · 30-60 days
3 · Human oversight degraded to rubber-stamp in practice · EU AI Act Art. 14, 26 · A · YELLOW · both · Redesign review workflow requiring substantive justification; introduce periodic challenge sessions · 30-60 days
4 · Technical documentation incomplete/outdated · EU AI Act Art. 11, Annex IV · G · YELLOW · operational · Complete Annex IV technical documentation covering architecture, data lineage, limitations · 30-60 days
5 · Audit trail only partially reconstructable · ISO/IEC 42001 Clause 8 · G · YELLOW · both · Implement version-controlled change log linked to named owner · 30-60 days
6 · CBUAE explainability/outsourcing controls partial · CBUAE AI/ML Guidance · G · YELLOW · both · Produce explainability pack; document vendor/outsourcing due diligence · 30-60 days
7 · Special-category data safeguards partial · GDPR Art. 9; UAE PDPL Art. 1, 10 · U · YELLOW · both · Map special-category data exposure; implement stronger safeguards where confirmed · 60-90 days
8 · Data minimisation/retention only partially configured · GDPR Art. 5; UAE PDPL Art. 5 · U · YELLOW · operational · Configure and document retention schedules and minimisation rules · 60-90 days
9 · Training/input data governance only partially documented · EU AI Act Art. 10 · U · YELLOW · policy · Document data governance record: source, collection, representativeness, update frequency · 60-90 days
10 · AI-literacy training limited to general awareness · EU AI Act Art. 4 · A · YELLOW · operational · Deliver CreditLens-specific training on limitations, failure modes, escalation triggers · 30-60 days
11 · Bias testing done once, no recurring schedule · EU AI Act Art. 10; DIFC Reg. 10.3 · R · YELLOW · both · Establish recurring bias-testing schedule with retained evidence · 30-90 days
12 · Incident response process improvised beyond kill-switch · EU AI Act Art. 26; NIST AI RMF Manage · R · YELLOW · policy · Document incident response plan: owner, rollback procedure, contact channel · 30-90 days
13 · Robustness/accuracy testing evidence thin · EU AI Act Art. 15 · R · YELLOW · operational · Commission structured accuracy/robustness/cybersecurity testing · 30-90 days
14 · AI disclosure inconsistent across touchpoints · EU AI Act Art. 50 · R · YELLOW · policy · Standardize disclosure language and placement across touchpoints · 30-90 days
15 · High-risk conformity assessment incomplete · EU AI Act Art. 16, 26, 43 · R · YELLOW · both · Complete in-progress conformity assessment with firm completion date · 30-90 days
16 · DIFC fairness/non-discrimination testing not systematic · DIFC Reg. 10.3 · R · YELLOW · operational · Systematize fairness testing into recurring schedule · 30-90 days
17 · Value tracking incomplete (only one of two dimensions) · NIST AI RMF Map 5 · D · YELLOW · operational · Document failure-cost/exposure estimate alongside value tracking · 0-30 days
18 · Decommissioning plan informal, not documented · ISO/IEC 42001 Clause 8.4 · D · YELLOW · policy · Draft written decommissioning plan · 0-30 days

Remediation Roadmap

Immediate (0-7 days)
- Begin drafting the failure-cost/exposure estimate to complete value tracking (Gap #17).
- Assign named ownership and kickoff date for the DPIA (Gap #1) and EU database registration (Gap #2).

Short-term (7-30 days)
- Complete decommissioning plan documentation (Gap #18).
- Complete value tracking framework (Gap #17).
- Standardize AI-disclosure language across touchpoints (Gap #14).

• Draft incident response plan documentation (Gap #12).

Medium-term (30-90 days)
• Complete DPIA (Gap #1).
• Complete EU database registration (Gap #2).
• Redesign human review workflow and deliver CreditLens-specific training (Gaps #3, #10).
• Complete Annex IV technical documentation and implement audit-trail change log (Gaps #4, #5).
• Produce CBUAE explainability pack and outsourcing documentation (Gap #6).
• Complete special-category data mapping and minimisation/retention configuration (Gaps #7, #8).
• Document training-data governance record (Gap #9).
• Establish recurring bias-testing schedule and systematize DIFC fairness testing (Gaps #11, #16).
• Commission robustness/accuracy/cybersecurity testing (Gap #13).
• Complete in-progress EU AI Act conformity assessment (Gap #15).

Ongoing (90+ days)
• Maintain recurring bias and fairness testing cadence (Gaps #11, #16) as a standing quarterly process.
• Maintain and refresh CreditLens-specific staff training on an annual cycle (Gap #10).
• Continue post-market monitoring and value-tracking review as standing governance functions.
• Extend Navigator assessment coverage to the advisory chatbot and CV-screening tool, and formalize a permitted-tools policy addressing informal shadow-AI use identified in the profile summary.

Next Steps

1. Complete the EU AI Act conformity assessment and EU database registration for CreditLens as the highest-priority binding regulatory obligations with direct enforcement exposure (Gaps #2, #15).
2. Commission a DPIA for CreditLens covering discriminatory-outcome risk before further processing of EU or UAE client data in lending decisions (Gap #1).
3. Address the human-oversight effectiveness gap directly — redesign the credit committee review process to eliminate rubber-stamping and deliver system-specific training (Gaps #3, #10).
4. Run separate GUARD Navigator assessments for the advisory chatbot and the third-party CV-screening tool, both of which carry independent high-risk/transparency obligations not captured in this report.
5. Formalize a permitted-tools policy and governance response to the disclosed informal use of general-purpose AI assistants by employees.

Disclaimer

This assessment assists with AI governance workflows and is produced by the GUARD GRC Assessment Engine based on responses provided through the GUARD Navigator questionnaire. It does not constitute legal advice. All findings, regulatory citations, and remediation recommendations should be reviewed by qualified legal professionals before being relied upon for compliance decisions. Regulatory instruments cited are based on the GUARD regulatory knowledge base and the provided GUARD database context as of 8 July 2026; organizations should verify currency of all cited provisions with qualified counsel.

Meridian Capital Partners

Company-wide · completed 2026-06-01 17:39:24 UTC · content SHA-256 c7fa488b634cc1bf…

AI System Assessment Report

System: Hybrid AI system (combination of multiple AI technologies) used for credit decisions, hiring/employment decisions, and customer service/chatbot functions
Client: Meridian Capital Partners
Assessment Date: 8 July 2026
Assessor: GUARD GRC
Jurisdictions: European Union/EEA, Dubai International Financial Centre (DIFC), United Arab Emirates (general)
Risk Classification: High-Risk (per stated use in credit decisions and employment/hiring — EU AI Act Article 6 high-risk triggers apply per user classification)

DATA QUALITY ADVISORY

Before proceeding, two data quality issues require attention:

1. Unverified supplementary profile data. An untrusted, unauthenticated "profile summary" was submitted alongside the Navigator JSON, containing claims not present in the structured intake (named systems "CreditLens" and a "CV-screening tool," DFSA regulatory status, specific Annex III sub-paragraph citations, headcount, and branch details). None of these claims are corroborated by actual questionnaire responses (no question in this intake asks about named AI systems, DFSA status, or branch locations). This assessment does not rely on that content. We strongly recommend Meridian resubmit these facts through the formal Navigator

intake (or a verified follow-up questionnaire) so they can be properly assessed and cited with regulatory basis. Treating unverified narrative claims as fact in a compliance report would misrepresent your actual regulatory exposure.

2. Contradiction in EU_06 answer. CLS_03 does not indicate Meridian operates a general-purpose AI model — the classification describes credit decisions, hiring, and customer service functions (task-specific applications). However, EU_06 was answered "Yes, and it has systemic risk potential," which describes a General-Purpose AI model under EU AI Act Article 51 (systemic-risk GPAI). This is inconsistent with a firm using AI for internal credit scoring, hiring, and a customer service chatbot — none of which are typically GPAI models with systemic risk designation. We recommend Meridian verify this answer. It is possible the chatbot uses a third-party GPAI model (e.g., a foundation model integrated via API), which would NOT make Meridian itself a GPAI provider subject to Article 51-55 obligations — those obligations fall on the model developer, not the deployer, in most cases. This assessment proceeds on the stated answer (EU_06_HIGH_IMPACT) but flags it as requiring clarification, as it materially affects which EU AI Act obligations apply.

Executive Summary

Meridian's AI system, as classified through the Navigator intake, is a hybrid AI solution used across three distinct high-stakes functions: credit/financial assessment decisions, hiring and employment decisions, and customer service automation. It processes personal data, financial data, and profiling data affecting individual citizens/consumers and employees/job applicants across EU/EEA, DIFC, and wider UAE jurisdictions. Given the credit-decision and hiring use cases, this system falls squarely within EU AI Act high-risk categories, and the multi-jurisdictional footprint (EU + DIFC + UAE) triggers overlapping data protection and AI governance obligations under GDPR, the DIFC Data Protection Law with Regulation 10, and (for UAE mainland touchpoints) the UAE PDPL.

The overall compliance posture is mixed to concerning: Governance (G) and Dollar Drain (D) are rated YELLOW, Attrition of Skills (A) is YELLOW trending toward RED-adjacent due to compounding gaps, Unauthorized Data Access (U) is YELLOW, and Reputation (R) is RED — driven by the absence of bias testing corroboration across two independent questions (DIFC_03_NO and GOV_10_PARTIAL), no incident response plan (GOV_11_NO), and only partial transparency to affected individuals (EU_04_PARTIAL, GOV_09_PARTIAL). This is a material finding given the system's direct use in credit and employment decisions — both classic high-risk, high-scrutiny domains under EU AI Act Article 6 and Annex III.

The Highest Common Denominator standard applicable here is the EU AI Act's high-risk obligations (Articles 9-15, 26, 50) combined with DIFC Regulation 10's AI-specific governance and fairness requirements — both regimes converge on requiring documented risk management, bias mitigation, human oversight, and transparency to affected persons. Meeting the EU AI Act high-risk bar will substantially satisfy DIFC Regulation 10 and vice versa, given their aligned principles (ethical design, fairness, transparency). This report identifies 13 findings requiring remediation (4 RED, 9 YELLOW). Additionally, several genuine compliance strengths were identified and are detailed in the category narratives below (these are not counted in the 13 findings).

Regulatory Applicability

Regulation · Applies · Trigger · Key Obligations

EU AI Act (Reg. 2024/1689) · Yes · CLS_01 includes EU/EEA; CLS_03 indicates credit decisions and hiring — both Annex III high-risk domains; EU_01_YES confirms high-risk self-assessment · Articles 9 (risk management), 14 (human oversight), 26 (deployer obligations), 50 (transparency)

GDPR (Reg. 2016/679) · Yes · CLS_01 includes EU/EEA; CLS_04 confirms personal data, financial data, and profiling data processed; affected parties include EU-based individuals · Articles 5-6 (lawfulness, DP_01 cites legitimate interest), 13 (transparency), 22 (automated decision-making, relevant given credit/hiring use), 35 (DPIA)

DIFC Data Protection Law (Law No. 5 of 2020) · Yes · CLS_01 includes DIFC · Parts 2-3 (lawful processing), Part 4 (cross-border transfer, Article 26)

DIFC Data Protection Regulation 10 · Yes · CLS_01 includes DIFC; AI system processes personal data via autonomous/semi-autonomous means (DIFC_01 through DIFC_04 answered, confirming DIFC AI-specific module activation) · Regulation 10.2 (transparency to data subjects), Regulation 10.3.1 (ethical, fairness, and bias-mitigation design requirements)

UAE Federal PDPL (Decree-Law 45/2021) · Conditional · CLS_01 includes UAE (general); applies where processing touches UAE mainland entities or data subjects outside the DIFC free zone · General data protection obligations parallel to DIFC law; verify scope of UAE mainland data flows

NIST AI RMF · Advisory · Multi-jurisdictional high-risk AI deployment; referenced as compliance benchmark in DIFC Regulation 10 guidance · GOVERN, MAP, MEASURE, MANAGE functions — used as structuring framework for remediation

ISO/IEC 42001:2023 · Advisory · Referenced in DIFC Regulation 10 guidance as an acceptable design/certification framework · AI management system controls, Annex A data governance provisions

Note on cross-border transfer regulation: DP_02 was answered "No, all data stays within the original jurisdiction." Based on this answer, GDPR Chapter V (Articles 44-49) transfer mechanisms and DIFC Data Protection Law

Article 26 transfer provisions are not currently triggered for cross-border movement, since no transfer was reported. This should be revisited if the unverified profile-summary claim of a Frankfurt branch sharing data with a Dubai HQ is confirmed through formal intake — that would constitute a DIFC-to-EU (or EU-to-DIFC) transfer requiring an adequacy basis or a transfer mechanism under Article 26/27 of the DIFC Data Protection Law and GDPR Chapter V.

GUARD Risk Assessment

G — Governance: YELLOW

Finding: Meridian has some governance foundations in place — leadership has visibility into AI risk (GOV_02_YES) and a named individual has been formally appointed as AI Senior Officer for DIFC operations (DIFC_01_YES), which is a genuine strength directly satisfying a DIFC Regulation 10 expectation. However, outside DIFC, AI accountability is informal (GOV_01_PARTIAL — someone handles it without formal appointment), the written AI policy is incomplete (GOV_04_PARTIAL), the AI system register is incomplete (DIFC_04_PARTIAL), DIFC risk classification work is only partially done (DIFC_02_PARTIAL), and — most critically — AI decisions cannot currently be traced back to who approved the system, what data it used, or why a specific decision was made (GOV_03_NO).

Regulatory basis: EU AI Act Article 26 (deployer obligations, including maintaining logs and instructions for use); DIFC Data Protection Regulation 10.2 (transparency and accountability for AI systems processing personal data); DIFC Regulation 10.3.1(a) (ethical design requiring traceable, unbiased decision-making); NIST AI RMF GOVERN function.

What's missing:
1. Formal, enterprise-wide (non-DIFC-specific) appointment of an accountable AI owner with documented authority (GOV_01_PARTIAL).
2. Complete and current written AI governance policy covering ethics, risk management, and compliance (GOV_04_PARTIAL).
3. Traceability/audit trail mechanism for AI decisions — no ability to reconstruct who approved a system, what data informed a decision, or why an outcome occurred (GOV_03_NO).
4. Comprehensive AI risk classification aligned to DIFC guidelines (DIFC_02_PARTIAL).
5. Complete, current AI system inventory/register (DIFC_04_PARTIAL).

Remediation:
1. Formally designate an accountable AI governance owner (e.g., extend the DIFC AI Senior Officer role group-wide, or appoint a parallel role for non-DIFC operations) with documented authority and escalation responsibilities.
2. Complete and approve a comprehensive written AI governance policy covering intended use, prohibited use, risk tiers, escalation, and review cadence.
3. Implement a decision-logging mechanism capturing model version, input data lineage, approval chain, and outcome rationale for each AI-driven decision, retained for audit purposes.
4. Complete the DIFC AI risk classification exercise for all systems, documenting risk tier and justification per system.
5. Complete and maintain a current AI system register recording purpose, owner, data types, and risk level for every AI system in use.

Effort: Medium
Timeline: 30-90 days

U — Unauthorized Data Access: YELLOW

Finding: Meridian relies on legitimate interest as its lawful basis for processing personal data in the AI system (DP_01_LEGIT) — this is a valid basis under GDPR Article 6(1)(f) and DIFC Data Protection Law Article 10, provided a documented legitimate interest assessment (balancing test) exists. No cross-border data transfer was reported (DP_02_NO), which is a positive finding reducing exposure under GDPR Chapter V and DIFC Law Article 26. However, a formal data protection risk assessment is still in progress rather than complete (DP_03_PROG), access controls to personal data are only partially restricted (DP_04_PARTIAL), and breach response procedures exist only partially (DP_05_PARTIAL).

Regulatory basis: GDPR Article 6(1)(f) (legitimate interest basis) and Article 35 (DPIA requirement, particularly relevant given automated decision-making in credit and hiring contexts under Article 22); DIFC Data Protection Law Article 10 (lawful processing) and Article 20 (data protection impact assessment); DIFC Regulation 10.2 (transparency requirements for AI processing).

What's missing:
1. Documented legitimate interest assessment (balancing test) has not been confirmed as complete — DP_01 confirms the basis but not that the assessment is documented.
2. Data protection/privacy risk assessment for the AI system is still in progress, not finalized (DP_03_PROG).

3. Access controls to personal data used by the AI system remain partially implemented, with identified gaps (DP_04_PARTIAL).

4. Data breach response plan exists only partially — notification procedures, damage mitigation, and regulatory reporting steps are incomplete (DP_05_PARTIAL).

Remediation:

1. Document a formal legitimate interest assessment (balancing test) evidencing that the business need does not override individual rights, and retain it as part of the Article 6 lawful basis record.

2. Complete the in-progress data protection risk assessment (DPIA) within a defined timeline, ensuring it specifically addresses the automated decision-making elements of the credit and hiring use cases (GDPR Article 22 considerations).

3. Close remaining access control gaps by implementing role-based access controls limiting personal data access strictly to personnel with a defined operational need, and conduct an access audit to identify current over-permissioning.

4. Finalize the data breach response plan, including specific notification timelines and templates for DIFC Commissioner reporting, GDPR supervisory authority reporting (72-hour rule under Article 33 equivalent obligations), and affected data subject communication procedures.

Effort: Medium
Timeline: 30-60 days

A — Attrition of Skills: YELLOW

Finding: Meridian operates a human-on-the-loop model (GOV_05_ONLOOP) — humans monitor and can intervene, which is an appropriate baseline oversight structure for high-risk AI use cases. However, this is significantly undermined by three compounding gaps: human review of AI outputs is inconsistent and depends on the reviewer and complexity ("sometimes" — GOV_06_PARTIAL, a genuine gap with mitigation, not full compliance), no formal training has been provided to people who use or oversee the AI system (GOV_07_NO), and workforce impact from AI replacing human tasks has not been formally assessed (GOV_08_NO). The combination of inconsistent critical review and zero formal training creates meaningful risk that human oversight is nominal rather than substantive — a critical concern for a human-on-the-loop model deployed in credit and hiring decisions.

Regulatory basis: EU AI Act Article 14 (human oversight — requiring that persons assigned oversight have "necessary competence, training and authority"); DIFC Regulation 10.3.1 (design principles requiring meaningful human accountability); NIST AI RMF GOVERN 1.3 (workforce competency for AI risk management).

What's missing:

1. No formal training program exists for personnel who use or oversee the AI system, undermining the effectiveness of the human-on-the-loop model (GOV_07_NO).

2. Human review of AI outputs is inconsistent — quality of scrutiny depends on the individual reviewer and complexity of the case, rather than being governed by a defined standard (GOV_06_PARTIAL).

3. No formal assessment of workforce impact has been conducted despite the AI system affecting hiring and credit decision workflows previously performed by humans (GOV_08_NO).

Remediation:

1. Design and deliver a mandatory training program for all personnel with AI oversight responsibilities, covering system capabilities, known limitations, common failure modes, and escalation triggers. Require completion before granting override/review authority.

2. Establish a documented human review standard defining minimum review criteria, required scrutiny level by decision type (e.g., mandatory secondary review for declined credit applications or rejected candidates), and quality-assurance spot-checks on review consistency.

3. Conduct a formal workforce impact assessment identifying which roles/tasks have been affected by AI deployment, and document support measures (redeployment, upskilling) for affected employees.

Effort: Medium
Timeline: 30-90 days

R — Reputation: RED

Navigator score: RED (3 gap indicators against 1 compliance indicator). Adjusted assessment: RED confirmed, but with one contextual override. EU_05_NO ("does not generate synthetic content") is reclassified as a POSITIVE/neutral compliance indicator per Rule 3 — it is a genuine fact reducing deepfake/synthetic-media risk, not a gap. This slightly improves the compliance indicator count (2 positives: EU_01_YES, EU_05_NO-reclassified) but does not change the overall severity, because the remaining substantive gaps — no bias testing corroborated across two independent questions, no incident response plan, and only partial transparency — are material and compounding in a high-risk credit/hiring context.

Finding: Meridian correctly self-identified its system as high-risk under the EU AI Act (EU_01_YES), which is a

positive indicator of self-awareness. The system does not generate synthetic content (EU_05_NO), removing deepfake-related transparency obligations under EU AI Act Article 50(4). However, disclosure to individuals interacting with the AI is only partial (EU_04_PARTIAL) — users are not comprehensively told they are dealing with AI or its limitations. Bias testing is a significant, corroborated gap: DIFC_03_NO confirms no bias testing has been performed at the DIFC-fairness-assessment level, while GOV_10_PARTIAL confirms only partial testing has occurred generally — these are not contradictory but describe the same underlying deficiency from two angles: bias testing is incomplete and does not meet a demonstrable fairness standard. Explainability to affected individuals is also only partial (GOV_09_PARTIAL). Most critically, there is no formal incident response plan for AI-caused harm, including remediation and disclosure procedures (GOV_11_NO) — a severe gap given the system's use in credit and hiring decisions where erroneous or discriminatory outcomes carry direct legal and reputational consequences. Public disclosure of AI use is also only partial (GOV_12_PARTIAL).

Regulatory basis: EU AI Act Article 50 (transparency obligations to natural persons interacting with AI systems); EU AI Act Article 26(5) (deployer monitoring and incident escalation obligations); DIFC Regulation 10.2.2 (transparency requiring sufficient detail for individuals to assess risk and object); DIFC Regulation 10.3.1(a)-(b) (ethical and fairness design principles, explicitly requiring bias mitigation); GDPR Article 22 (right to explanation/ human intervention for automated decisions affecting credit and employment).

What's missing:

1. No comprehensive AI bias testing has been completed — confirmed as absent at the DIFC fairness-assessment level (DIFC_03_NO) and only partially addressed generally (GOV_10_PARTIAL), a critical gap given use in credit and hiring decisions where discriminatory outcomes carry direct legal exposure under equality/anti-discrimination law.
2. No formal incident response plan exists for AI-caused harm, including remediation and disclosure procedures (GOV_11_NO).
3. AI disclosure to individuals interacting with the system is incomplete — not comprehensive regarding what the system can and cannot do (EU_04_PARTIAL).
4. Explanation of AI-driven decisions to affected individuals is only partially available (GOV_09_PARTIAL).
5. Public disclosure of AI use (transparency reporting, website disclosure) is limited, largely confined to privacy notices (GOV_12_PARTIAL).

Remediation:

1. Conduct a comprehensive, documented bias assessment across all protected characteristics (age, gender, race, disability, etc.) for the credit-decision and hiring components of the system, using a structured statistical fairness testing methodology, and remediate any identified disparities before continued high-risk deployment.
2. Draft and implement a formal AI incident response plan covering detection, containment, remediation, affected-party notification, and regulatory disclosure procedures, aligned to EU AI Act Article 26(5) serious incident reporting obligations and DIFC Commissioner notification expectations.
3. Enhance AI disclosure to end users to comprehensively state that they are interacting with an AI system, its intended purpose, and its known limitations, delivered at first interaction point (chatbot/credit/hiring touchpoints).
4. Build a structured explainability mechanism enabling affected individuals to request and receive a meaningful explanation of how and why an AI-driven decision was made, satisfying GDPR Article 22(3) and DIFC Regulation 10.2.2 transparency requirements.
5. Expand public AI disclosure beyond privacy notices — publish a plain-language AI use statement describing where and how AI is used in decision-making that affects customers and employees.

Effort: High
Timeline: 0-60 days (bias testing and incident response plan are immediate priorities; disclosure enhancements follow)

D — Dollar Drain: YELLOW

Finding: Meridian actively monitors whether the AI system delivers business value (GOV_13_YES), a genuine strength indicating value governance discipline. However, financial exposure from potential non-compliance or legal liability has only been partially estimated (GOV_14_PARTIAL), systematic testing and validation of AI performance — including accuracy, fairness, safety, and drift detection — is not conducted (GOV_15_NO), and the AI-specific risk management framework is incomplete (GOV_16_PARTIAL). The absence of systematic performance testing (GOV_15_NO) compounds the Reputation-category bias testing gap and increases financial exposure, since undetected model drift or degraded accuracy in credit decisions directly creates regulatory and litigation risk.

Regulatory basis: EU AI Act Article 99 (penalty structure for high-risk AI Act non-compliance, including obligations under Articles 9 and 14 relevant here); GDPR Article 83 (administrative fines framework, relevant given the automated decision-making and profiling data processed); DIFC Data Protection Law enforcement provisions (Commissioner powers under the Law).

What's missing:

1. No comprehensive financial exposure estimate exists for potential AI-related non-compliance or legal liability

— only partial cost identification has occurred (GOV_14_PARTIAL).
2. No systematic, ongoing testing and validation regime exists to check AI accuracy, fairness, safety, and performance drift (GOV_15_NO).
3. The AI-specific risk management framework is incomplete/immature (GOV_16_PARTIAL).

Remediation:
1. Conduct a comprehensive financial exposure assessment quantifying potential regulatory penalty exposure (under EU AI Act Article 99 and GDPR Article 83 fine tiers), remediation costs, and professional liability risk specific to the credit and hiring AI use cases.
2. Implement a systematic, recurring testing and validation program for the AI system covering accuracy benchmarking, fairness metrics, safety checks, and performance drift detection, with defined thresholds triggering escalation or model retraining.
3. Complete and formalize the AI-specific risk management framework, ensuring it is distinct from general IT security frameworks and specifically addresses AI lifecycle risks (data, model, deployment, monitoring).

Effort: Medium
Timeline: 30-90 days

Clarification Required

· Question · User Response · Issue · Impact on Assessment
1 · EU_06 — Is your AI system a general-purpose AI model with systemic risk potential? · "Yes, and it has systemic risk potential" · Inconsistent with CLS_03 (credit decisions, hiring, customer service — task-specific use cases), which do not typically describe an organization operating a systemic-risk GPAI model. May reflect confusion about whether a third-party foundation model embedded in the chatbot makes Meridian a GPAI provider versus deployer. · If confirmed as stated, Meridian would face EU AI Act Article 51-55 GPAI systemic-risk obligations (model evaluation, adversarial testing, incident reporting, cybersecurity) in addition to high-risk deployer obligations — a materially larger compliance scope. If this is a misclassification and Meridian is merely a deployer of a third-party GPAI model, Article 51-55 obligations fall on the model provider, not Meridian. This must be verified before finalizing EU AI Act scope.
2 · Untrusted profile-summary attachment (not a Navigator question) · Claims regarding named systems, DFSA status, branch structure, headcount · Content was submitted through an unauthenticated channel outside the structured Navigator intake and contains no corroborating question/answer trail · Not incorporated into findings above. Recommend Meridian formally submit these details through a verified intake or documentation review process so they can be properly assessed and cited.

Compliance Gap Summary

· Finding · Regulation · GUARD Category · Severity · Remediation Type · Remediation · Timeline
1 · No bias testing performed/incomplete for credit and hiring decisions · DIFC Regulation 10.3.1(a)-(b); EU AI Act Art. 10 · R · RED · operational · Conduct comprehensive statistical bias assessment across protected characteristics; remediate disparities · 0-30 days
2 · No formal AI incident response plan · EU AI Act Art. 26(5); DIFC Data Protection Law · R · RED · both · Draft and implement incident response plan with detection, notification, and regulatory disclosure procedures · 0-30 days
3 · AI decisions cannot be traced (approval, data, rationale, accountability) · EU AI Act Art. 26; DIFC Regulation 10.3.1(a) · G · RED · both · Implement decision-logging mechanism capturing model version, data lineage, approval chain, rationale · 30-60 days
4 · No systematic AI performance/drift testing · EU AI Act Art. 99 (non-compliance exposure); GDPR Art. 35 · D · RED · operational · Implement recurring testing program for accuracy, fairness, safety, and drift detection · 30-60 days
5 · Inconsistent human review quality of AI outputs · EU AI Act Art. 14 · A · YELLOW · both · Establish documented human review standard with mandatory scrutiny criteria and QA spot-checks · 30-60 days
6 · No formal training for AI oversight personnel · EU AI Act Art. 14 · A · YELLOW · both · Design and deliver mandatory AI training program for oversight personnel · 30-60 days
7 · No workforce impact assessment for AI-displaced tasks · EU AI Act recital context; DIFC Regulation 10 · A · YELLOW · operational · Conduct formal workforce impact assessment and document support measures · 30-90 days
8 · Incomplete AI governance policy · DIFC Regulation 10.2; NIST AI RMF GOVERN · G · YELLOW · policy · Complete and approve comprehensive written AI governance policy · 30-60 days
9 · Incomplete AI risk classification and system register · DIFC Regulation 10 · G · YELLOW · both · Complete DIFC risk classification exercise and finalize AI system register · 30-90 days
10 · Partial data protection risk assessment (in progress) · GDPR Art. 35; DIFC Data Protection Law Art. 20 · U · YELLOW · operational · Complete DPIA addressing automated decision-making elements · 30-60 days
11 · Partial access controls to personal data · GDPR Art. 5, 32 (security); DIFC Data Protection Law · U · YELLOW · operational · Implement role-based access controls and conduct access audit · 0-30 days
12 · Partial data breach response plan · GDPR Art. 33; DIFC Data Protection Law · U · YELLOW · both · Finalize breach response plan with notification timelines and templates · 0-30 days
13 · Incomplete financial exposure estimate and AI risk management framework · EU AI Act Art. 99; GDPR Art.

83 · D · YELLOW · both · Conduct financial exposure assessment; complete AI-specific risk management framework · 30-90 days

Remediation Roadmap

Immediate (0-7 days)
- Begin bias assessment scoping for credit and hiring AI functions (Gap #1).
- Begin drafting incident response plan structure (Gap #2).
- Initiate access control audit to identify over-permissioned accounts (Gap #11).
- Clarify EU_06 classification internally (systemic-risk GPAI vs. third-party model deployer) — see Clarification Required #1.

Short-term (7-30 days)
- Complete bias assessment and begin remediation of identified disparities (Gap #1).
- Finalize and approve incident response plan (Gap #2).
- Complete access control remediation (Gap #11).
- Finalize data breach response plan with notification templates (Gap #12).

Medium-term (30-90 days)
- Implement AI decision-logging/traceability mechanism (Gap #3).
- Deploy systematic performance/drift testing program (Gap #4).
- Establish human review standard and deliver AI oversight training (Gaps #5, #6).
- Conduct workforce impact assessment (Gap #7).
- Complete AI governance policy (Gap #8).
- Complete DIFC risk classification and AI system register (Gap #9).
- Complete DPIA (Gap #10).
- Conduct financial exposure assessment and complete AI risk management framework (Gap #13).

Ongoing (90+ days)
- Recurring bias and performance testing cadence (building on Gap #1, #4).
- Periodic review and refresh of AI governance policy, risk classification, and system register (building on Gaps #8, #9).
- Continued value-monitoring practice (existing strength — GOV_13_YES) extended to include compliance ROI tracking.
- Periodic public AI disclosure updates and transparency reporting maturity.

Next Steps

1. Verify the EU_06 classification internally with technical/product teams to determine whether Meridian is a GPAI provider with systemic risk (Article 51-55 obligations) or a deployer of a third-party foundation model — this materially changes EU AI Act scope and may warrant a formal legal opinion on Annex III/GPAI classification if ambiguity remains after internal review.
2. Prioritize the bias assessment and incident response plan (Gaps #1 and #2) given the RED rating in Reputation and the direct use of AI in credit and hiring decisions — these are the highest enforcement-exposure items.
3. Resubmit organizational facts through a verified channel (not the untrusted profile-summary attachment) so that named systems, regulatory registrations, and operational structure can be properly assessed and cited in a follow-up review.
4. Implement the decision-traceability mechanism (Gap #3) as a foundational control — most other governance and reputation gaps become easier to close once decisions are logged and auditable.
5. Complete the in-progress DPIA and finalize the AI governance policy within the next 30-60 days to close the two most foundational documentation gaps underpinning multiple other findings.

Disclaimer

Appendix B — Attestations (full text)

AI Transparency Statement — Meridian Capital Partners

AI Transparency Statement — Meridian Capital Partners

§ 1. About this statement

Meridian Capital Partners ("Meridian", "we", "us") is an investment advisory and wealth management firm headquartered in Dubai (DIFC) with a branch office in Frankfurt serving EU-domiciled clients. We use artificial intelligence ("AI") in parts of our business, and we are publishing this statement so that clients, prospective clients, job applicants, and other individuals who interact with us understand what those AI systems do and how we govern them.

This statement addresses the AI-specific transparency obligations that apply to us given our operating footprint: DIFC Regulation 10 (including Regulation 10.2.2, on transparency for automated processing of personal data), the EU AI Act (Article 50 transparency obligations for providers and deployers), GDPR (Article 22, automated individual decision-making), and the UAE Federal PDPL (automated-processing provisions, to the extent our processing touches UAE mainland data subjects). We have not referenced regimes or obligations beyond these, as our inputs do not evidence a broader footprint.

This statement is a companion to our Privacy Notice. For lawful bases, retention periods, and the full catalogue of data-protection rights, please refer to the Privacy Notice; this document focuses specifically on our AI systems.

§ 2. Our approach to responsible AI

We maintain an internal register of the AI systems we operate, which records, for each system, its purpose, risk classification, the categories of data it uses, and its applicability against relevant regimes (including the EU AI Act, GDPR, DIFC Regulation 10, UAE PDPL, and CBUAE model-risk expectations). Systems in that register are assessed against these regimes, and gaps identified through that assessment are tracked for remediation.

We are being direct about the current state of that governance work, consistent with our commitment to transparency:

- A named individual holds AI governance responsibility for our DIFC operations. Formal, enterprise-wide appointment of an accountable AI owner covering all jurisdictions is [to be completed by the organisation — a group-wide AI governance role has not yet been evidenced].
- Our AI system register and risk classification exercise are partially complete and are being finalised.
- A documented, enterprise-wide AI governance policy is in progress, not yet finalised.
- Decision traceability (the ability to reconstruct who approved a system, what data informed a decision, and why an outcome occurred) is not yet fully in place for our credit scoring system; this is a priority remediation area.
- Bias/fairness testing for systems affecting credit and hiring outcomes is incomplete and is being addressed as a priority.
- A formal AI incident response plan (covering detection, remediation, and notification for AI-related harm) does not yet exist and is being developed.

We do not claim a mature AI governance programme beyond what is evidenced above. We disclose these gaps because we believe you are entitled to know the true state of oversight over systems that may affect you, not only the systems themselves.

Separately, we are aware that staff may use general-purpose AI assistants informally, without a defined organisational policy governing permitted tools. This is an identified gap we are working to close through a formal acceptable-use policy; it does not currently affect the systems described in § 3, which are formally deployed and monitored.

§ 3. The AI systems we operate

3.1 CreditLens (credit and counterparty scoring)

CreditLens is an in-house model that scores clients for creditworthiness in connection with lending, margin, and counterparty-limit decisions. It is built by us (we are the "provider" for EU AI Act purposes) using internal repayment history, credit bureau data, and transaction features. It does not interact directly with individuals — it is a back-office scoring tool that feeds our credit committee's workflow. You are not "chatting" with CreditLens; rather, if you apply for or hold a credit facility or margin arrangement with us, its output may inform decisions about your facility or limit, including for clients of our Frankfurt branch.

Data used: personal and financial data (repayment history, bureau data, transaction features).

Content generation: CreditLens produces a score/output for internal use; it does not generate content presented

directly to you.

3.2 Client Advisory Chatbot

The Advisory Chatbot is a conversational assistant available on our retail digital advisory channel (web portal and mobile app), built on a third-party large language model (GPT-4o, via API). It answers product and portfolio questions and provides general market commentary. It is designed to route any query involving actual investment advice to a licensed human adviser — it is not intended to give regulated advice itself.

Because you interact with this system directly, you are told that you are dealing with an AI assistant rather than a human when using the chat channel, consistent with our transparency obligations. All conversations are logged.

Data used: the personal and account/portfolio data needed to answer your query.

Content generation: yes — the chatbot generates conversational responses and market commentary in real time. It does not generate images, audio, or video, and does not produce synthetic media.

3.3 TalentScreen CV Ranker

TalentScreen is a third-party tool used by our HR function to rank inbound job applications against role profiles, currently used for front-office and technology hiring and running as a pilot. It processes candidate employment history and education data to produce a shortlist ranking. It does not make final hiring or rejection decisions on its own — a recruiter reviews rankings before any rejection communication is sent.

Job applicants do not interact with this tool directly; it operates behind the scenes on submitted application materials.

Data used: candidate personal data, including employment history and education.

Content generation: no — it ranks/scores existing candidate information; it does not generate new content.

§ 4. Automated decision-making

CreditLens materially supports decisions with a legal or similarly significant effect on individuals — namely, decisions to decline a credit facility or reduce a lending/margin limit. This can affect retail and SME clients, including EU-domiciled clients of our Frankfurt branch.

- Nature of the decision: approval, decline, or limit-adjustment of lending or margin facilities.
- Human involvement: scores feed into a credit committee workflow; the intended model is human review of scores rather than fully automated decline. However, we cannot currently demonstrate, in all cases, that a human meaningfully reviews and can override every adverse outcome before it takes effect — our ability to trace who approved a given decision and why is an identified gap (see § 2). We are treating this as a priority area for remediation, and until it is closed, we recommend that affected clients exercise the review rights in § 6.
- Safeguards: the score is one input into a broader credit committee process; a documented bias/fairness assessment of this model has not yet been completed and is in progress.

TalentScreen ranks candidates but does not itself reject or hire — a recruiter reviews shortlists before any rejection email is sent. We do not treat this as a decision made "solely" by automated means, but we recognise the ranking materially influences which candidates progress, so the safeguards in § 5 apply.

The Advisory Chatbot does not make decisions with legal or similarly significant effect — it answers questions and routes advice-related queries to a human adviser. No automated investment decision is made by the chatbot itself.

Beyond these, we are not aware of other systems in our register that make or materially support decisions of this kind.

§ 5. Human oversight

- CreditLens operates on a human-on-the-loop basis: humans (the credit committee) are intended to monitor and review scores as part of the lending decision process. However, the consistency of that review, and the ability to trace and reverse a given output, is not yet fully evidenced or standardised — this is an identified gap we are working to close, including through mandatory training for reviewers and a defined review standard.
- TalentScreen shortlists are reviewed by human recruiters before any rejection communication is issued, giving recruiters the ability to override the tool's ranking at that stage.
- Advisory Chatbot conversations are logged and monitored; any query identified as seeking investment advice is escalated to a licensed human adviser rather than being answered by the system.

We have not evidenced a formal, enterprise-wide incident response plan for AI-related harm, or systematic ongoing performance/drift testing of these systems. Both are being developed. [To be completed by the organisation: confirmation of testing cadence and incident response procedure once finalised.]

§ 6. Your rights and choices

Depending on which system and jurisdiction applies to you, you may have the following AI-specific rights. This list is a summary — for the full data-protection rights catalogue, see our Privacy Notice.

• To be told you are interacting with an AI system. Applies to the Advisory Chatbot (EU AI Act Article 50; DIFC Regulation 10.2.2).

• To obtain human intervention / request human review of a significant automated decision, and to express your point of view. Applies principally to CreditLens outcomes affecting lending or margin facilities (GDPR Article 22; DIFC Data Protection Law Article 38; UAE PDPL Article 18).

• To contest an automated or AI-supported decision that produces legal or similarly significant effects concerning you (GDPR Article 22; DIFC Data Protection Law Article 38; UAE PDPL Article 18).

• To object to processing carried out by automated means, subject to the exceptions set out in the applicable regime (e.g., contractual necessity, consent, legal authorisation) (GDPR Article 22(2); DIFC Data Protection Law Article 38(2); UAE PDPL Article 18(2)).

• To request meaningful information about the logic of an AI-supported decision affecting you, to the extent this is not already provided in the Privacy Notice.

To exercise any of these rights, please use the contact details in § 8. We will handle your request in line with the process described in our Privacy Notice; where a request concerns a decision made using CreditLens or TalentScreen, it will be directed to the relevant business owner for human review.

§ 7. How we keep this statement accurate

We intend to review this statement at least annually, with the next scheduled review by 2027-07-09. We will also update it sooner where:

• a new AI system is added to our AI systems register that affects individuals;

• an existing system is reclassified (for example, a change in risk class or regulatory applicability);

• a material change is made to how a system works, what data it uses, or the degree of human oversight applied to it; or

• applicable law changes in a way that affects the disclosures in this statement.

The current version of this statement is available on our website alongside our Privacy Notice. [To be completed by the organisation: confirm publication URL.]

§ 8. Contact

If you have questions about our use of AI, or wish to exercise any of the rights described in § 6, please contact us at:

[To be completed by the organisation — no Data Protection Officer or dedicated AI contact is named in the inputs provided; please insert the appropriate contact point, e.g., DPO email/postal address.]

Statement date: 2026-07-09
Next review due: 2027-07-09

Privacy Notice — Meridian Capital Partners

Privacy Notice — Meridian Capital Partners

Notice date: 9 July 2026 | Next scheduled review: 9 July 2027

§ 1. Who we are

Meridian Capital Partners ("Meridian", "we", "us") is an investment advisory and wealth management firm registered in the Dubai International Financial Centre (DIFC), with a branch office in Frankfurt, Germany serving EU-domiciled clients. For the processing described in this notice, we act as the controller of your personal data.

Our full registered address, company registration number, and Frankfurt branch registration details are [To be completed by the organisation].

Contact us: [To be completed by the organisation — a general privacy contact (email and/or postal address) covering both our Dubai headquarters and Frankfurt branch].

Data Protection Officer: No DPO has been identified in our records at the time of this notice. [To be completed by the organisation: name and contact details of a DPO, or confirmation that no DPO is legally required together with identification of an accountable privacy contact].

§ 2. Scope of this notice

This notice applies to:

• Prospective and existing clients — retail and institutional — who receive investment advisory, wealth management, discretionary mandate, or lending/margin services from us, including through our digital advisory channel;

• Website and digital channel users, including visitors who interact with our client-facing advisory chatbot; and

• Job applicants whose CVs are processed through our recruitment screening tool.

This notice covers our processing of personal data in connection with investment advisory, wealth management, lending and margin, digital advisory, and recruitment activities across our DIFC and EU (Frankfurt) operations. It does not address our processing of current employees' data in the employment relationship itself, which — if applicable — is the subject of a separate notice.

§ 3. Personal data we process and why

The personal data we process depends on your relationship with us:

• Identification and contact data (name, contact details, identification documents) — to establish and administer client relationships.
• Financial data (income, assets, transaction history, portfolio holdings, repayment history, credit bureau data) — to provide advisory, discretionary mandate, and lending/margin services, and to assess creditworthiness, including through our credit scoring model.
• Communications data (chatbot conversation logs, correspondence, and, where applicable, call recordings) — to deliver and supervise our services and meet regulatory record-keeping requirements.
• Employment history and education data (job applicants) — to assess candidates against role requirements as part of recruitment.

We have not identified any current processing of special category or sensitive data (such as health, biometric, or criminal-records data) within our advisory, lending, or recruitment systems. If this changes, we will update this notice accordingly.

Purposes and lawful bases:

Purpose · GDPR basis (EU clients / Frankfurt branch) · DIFC Data Protection Law / UAE PDPL basis
Providing advisory and discretionary mandate services · Performance of a contract (Art. 6(1)(b)) · Necessary for performance of a contract
Credit, lending and margin decisions (including automated scoring) · Performance of a contract; legitimate interests in prudent risk management (Art. 6(1)(f)) · Legitimate interests / contractual necessity
Operating the client advisory chatbot · Legitimate interests in efficient client service (Art. 6(1)(f)) · Legitimate interests
Recruitment and CV screening · Steps taken prior to entering a contract (Art. 6(1)(b)); legitimate interests · Legitimate interests / pre-contractual steps
Regulatory compliance (financial services conduct, AML, record-keeping) · Legal obligation (Art. 6(1)(c)) · Compliance with an applicable legal obligation

Where we rely on legitimate interests, we consider whether those interests are overridden by your rights and interests. A documented record of that balancing assessment is [To be completed by the organisation].

§ 4. AI and automated decision-making

We use AI systems in parts of our business:

• CreditLens — an in-house credit and counterparty scoring model used to support lending, margin, and counterparty-limit decisions.
• Client Advisory Chatbot — a conversational assistant on our digital advisory channel that answers product and portfolio questions and routes advice requests to a human adviser.
• TalentScreen CV Ranker — a third-party tool used in recruitment to rank candidate CVs against role profiles.

Automated decision-making: CreditLens can influence outcomes with a legal or similarly significant effect on you, such as a declined facility or reduced credit limit. Where such a decision is based solely on automated processing, you have the right to request human review, to express your point of view, and to contest the decision. The advisory chatbot does not make decisions with legal or similarly significant effect and always routes advice requests to a licensed human adviser. CV screening outcomes are reviewed by a recruiter before any rejection is communicated, so this system does not currently produce solely automated decisions of that kind; this will be reviewed if the process changes.

Your rights in relation to AI-driven decisions: regardless of which system is involved, if you are ever subject to a decision based solely on automated processing that produces a legal or similarly significant effect on you, you may request human intervention, obtain an explanation, express your views, and contest the decision. We are also strengthening bias testing, decision traceability, and explainability across these systems as part of an ongoing governance programme.

For detailed, per-system disclosures — including data used, risk classification, human-oversight arrangements, and safeguards — please see our separate AI Transparency Statement, available on request or via our website.

§ 5. Data sharing and recipients

We may share your personal data with:

• OpenAI, as the provider of the underlying generative AI model powering the client advisory chatbot;

• TalentScreen FZ-LLC, as the provider of the CV-screening tool used in recruitment;
• Regulators and supervisory authorities, where required by applicable financial services or data protection law; and
• Professional advisers and service providers who support our operations (for example, IT infrastructure and compliance support). Specific categories of recipients beyond those named above are [To be completed by the organisation].

We do not sell your personal data.

§ 6. International transfers

We operate from the DIFC (Dubai) with a branch in Frankfurt, Germany, and serve both UAE-based and EU-domiciled clients. This footprint may involve transfers of personal data between our DIFC and EU operations, and to third-party AI vendors that may process data outside your home jurisdiction.

The specific safeguard mechanism(s) we rely on for any such transfers (for example, an adequacy finding, a DIFC-approved transfer mechanism, or EU Standard Contractual Clauses) are [To be completed by the organisation]. We will update this notice once our transfer mapping and safeguard documentation is finalised.

§ 7. How long we keep your data

We retain personal data for as long as necessary to provide our services, meet our regulatory and legal obligations (including financial services record-keeping requirements), and to establish, exercise, or defend legal claims.

Specific retention periods for each category of data (for example, client records, chatbot conversation logs, unsuccessful candidate CVs, and credit scoring inputs) are [To be completed by the organisation]. Once defined, we will publish these periods, or the criteria used to determine them, here.

§ 8. Your rights and how to exercise them

Depending on your location and the law that applies to your data (GDPR for EU-connected processing; the DIFC Data Protection Law for DIFC-connected processing; the UAE Federal PDPL where it applies), you may have some or all of the following rights:

• Access — to obtain confirmation of, and a copy of, the personal data we hold about you.
• Rectification — to have inaccurate or incomplete data corrected.
• Erasure — to request deletion of your data, subject to legal and regulatory retention requirements.
• Restriction of processing — to limit how we use your data in certain circumstances.
• Objection — to object to processing based on our legitimate interests.
• Data portability — available under GDPR, and in equivalent circumstances under DIFC law, where processing is based on consent or contract and carried out by automated means.
• Withdraw consent — where processing is based on consent, at any time, without affecting the lawfulness of processing before withdrawal.
• Rights relating to automated decision-making — to obtain human intervention, express your point of view, and contest a decision produced solely by automated means that has a legal or similarly significant effect on you (see § 4).

To exercise any of these rights, contact us using the details in § 10. We may need to verify your identity before responding, and some rights may be subject to exemptions or limitations under applicable law.

§ 9. How we protect your data

We apply organisational and technical measures designed to protect your personal data against unauthorised access, loss, misuse, or disclosure, including access controls over the systems described in this notice. We are actively strengthening our access controls, breach-response procedures, and AI-specific risk management as part of an ongoing governance programme. Details of specific technical security measures or vendor security certifications are [To be completed by the organisation].

§ 10. Contact, complaints and updates

To contact us about this notice, or to exercise your rights, use the contact details in § 1.

To complain, you may first raise concerns directly with us. You also have the right to lodge a complaint with the supervisory authority relevant to your location:

• DIFC-connected processing: the DIFC Commissioner of Data Protection.
• EU-connected processing (Frankfurt branch and EU clients): the competent EU/German data protection supervisory authority for the Frankfurt branch, or your own local EU supervisory authority.
• UAE mainland-connected processing, where the UAE Federal PDPL applies: the UAE Data Office.

Updates to this notice: We may update this notice from time to time, including as we complete the governance

work referenced above (marked [To be completed by the organisation]). We will notify material changes through our website and, where appropriate, directly to affected individuals. This notice was issued on 9 July 2026 and is scheduled for review by 9 July 2027.

Data Protection Impact Assessment: CreditLens Credit and Counterparty Scoring Model

Data Protection Impact Assessment: CreditLens Credit and Counterparty Scoring Model

§ 1. Identification of the processing activity

1.1 Controller identity

Meridian Capital Partners, a DIFC-registered entity headquartered in Dubai with a branch office in Frankfurt serving EU-domiciled clients. No further entity-type suffix, free-zone designation beyond DIFC, or registered address is stated in the inputs. [To be completed by the controller] — full legal entity registration number, registered address, and DFSA licence reference are required to complete controller identification.

1.2 System description

CreditLens (version 4.1) is an in-house, gradient-boosted credit and counterparty scoring model developed and operated by Meridian as Provider. It is trained on internal repayment history, bureau data, and transaction features. The model feeds the credit committee workflow and produces scores that can lead to declined facilities or reduced limits for retail and SME clients, including EU-domiciled clients of the Frankfurt branch. Lifecycle state: production. Autonomy level: human-on-the-loop. User impact level: high.

1.3 Purpose(s) of processing

The stated purpose is credit and counterparty scoring to support lending, margin, and counterparty-limit decisions within a regulated financial-services context (investment advisory, wealth management, and lending). The assessment session confirms CreditLens outputs have legal or similarly significant effects on data subjects (declined facilities, reduced limits).

1.4 Categories of personal data processed

Data categories processed: personal and financial (confirmed by CLS_04 in the assessment session, which notes the "no personal data" carve-out does not apply). Sources include internal repayment history, bureau data, and transaction features. Whether special-category data (e.g., health data referenced in insurance/lending contexts, or proxies correlating with protected characteristics) enters the model is unconfirmed — the assessment session records this as DPS_03_PART (special-category safeguards only partially implemented) and flags it as a data-mapping requirement. [To be completed by the controller] — a completed special-category data mapping exercise is required.

1.5 Categories of data subjects

Retail and SME clients subject to lending, margin, or counterparty-limit decisions, including EU-domiciled clients served from the Frankfurt branch and UAE-based clients served from the Dubai headquarters. No children's data is referenced in the inputs.

1.6 Recipients and processors

No named processors, sub-processors, or vendor/outsourcing arrangements for CreditLens components are documented in the inputs, although the assessment session identifies "CBUAE explainability/outsourcing controls partial" as an open gap requiring documentation of vendor due diligence for third-party components feeding the model. [To be completed by the controller] — identity of any third-party data or model-component providers, and applicable data processing agreements.

1.7 Cross-border data transfers

The assessment session states that Meridian "has mapped and safeguarded cross-border transfers of this system's data" (DPS_04_YES), reflecting DIFC-to-EU-to-UAE data flows implied by the Frankfurt branch structure. However, the specific transfer mechanism (e.g., EU Commission adequacy decision, GDPR Chapter V Standard Contractual Clauses, DIFC Data Protection Law Article 26 adequacy or Article 27 safeguards) is not stated in the inputs. [To be completed by the controller] — the specific legal transfer mechanism(s) relied upon for DIFC!”EU!”UAE flows.

1.8 Lawful basis

The assessment session records that Meridian "has a documented lawful basis for processing" (DPS_01_YES) as a genuine compliance strength, but the specific Article 6 GDPR lawful basis (e.g., contract, legitimate interests) and the corresponding UAE PDPL Article 6 basis are not specified in the inputs. [To be completed by the

controller] — the specific lawful basis relied upon per jurisdiction, and, if legitimate interests is relied upon, the balancing test documentation.

§ 2. Necessity and proportionality assessment

2.1 Necessity of the processing

Credit and counterparty scoring is inherent to the firm's lending and margin-decision function within its regulated financial-services business. The processing of repayment history, bureau data, and transaction features is directly relevant to assessing creditworthiness for lending decisions. However, the assessment session identifies that training/input data governance — "relevance, representativeness, and provenance" — is only partially documented (EUS_05_PART), meaning the necessity of the specific data fields used cannot yet be fully verified against the stated purpose.

2.2 Proportionality of the processing

The system operates in human-on-the-loop mode, which is intended to provide a proportionality safeguard against fully automated adverse decisions. However, the assessment session's most serious finding is that this human oversight has "become largely a rubber stamp in practice" (GOV_04_NO) — meaning the formal proportionality safeguard is not operating as designed. This materially undermines the proportionality case for legal-effect automated scoring absent effective human review.

2.3 Data minimisation

Data minimisation and retention limits on inputs, prompts, and logs are "only partially configured" (DPS_05_PART per the assessment session). No retention schedule is documented in the inputs. [To be completed by the controller] — documented retention schedules and minimisation rules for all fields, prompts, and logs used by CreditLens.

2.4 Alternative, less intrusive means

No assessment of alternative, less data-intensive scoring approaches (e.g., reduced feature sets, alternative model architectures with lower discriminatory-outcome risk) is documented in the inputs. [To be completed by the controller] — analysis of whether the current feature set and model architecture represent the least intrusive means of achieving the credit-scoring purpose, particularly in light of the open bias-testing and data-governance gaps identified below.

§ 3. Risks to data-subject rights

3.1 Risk: Absence of a completed DPIA prior to processing

Description: CreditLens produces automated scores with legal or similarly significant effects (declined facilities, reduced limits) on data subjects, yet no DPIA has been completed despite this being a likely mandatory trigger.

Likelihood: High — the assessment session confirms directly that "No DPIA has been conducted for a system that scores individuals and produces legal/significant effects" (DPS_02_NO), and this DPIA document itself is being produced retrospectively to close that gap.

Severity: High — absence of a DPIA for a high-risk, legal-effect automated decision-making system means risks to data subjects (including discriminatory outcomes) have not been systematically assessed prior to processing, contrary to the "prior" requirement in GDPR Art. 35(1) and UAE PDPL Art. 21(1).

Affected rights: Right to protection of personal data (GDPR Art. 35); right not to be subject to a decision based solely on automated processing producing legal effects (GDPR Art. 22); UAE PDPL Art. 21 impact-assessment right; DIFC Data Protection Law Article 20 (data protection impact assessment).

Evidence basis: Gap — "No DPIA completed despite likely mandatory trigger" | status: open | basis: GDPR Art. 35; UAE PDPL Art. 21 | inherent risk: yellow.

3.2 Risk: Degraded human oversight (automation bias)

Description: Human review of CreditLens outputs has, per the controller's own disclosure, become "largely a rubber stamp in practice," meaning credit committee reviewers are not meaningfully exercising independent judgment over adverse scoring outcomes.

Likelihood: High — this is a directly reported, current-state finding (GOV_04_NO), not a hypothetical risk.

Severity: High — where oversight has degraded to a rubber stamp, data subjects effectively receive fully automated decisions with legal effect (declined facilities, reduced limits) without the safeguard GDPR Art. 22 and EU AI Act Art. 14 require, despite the system nominally being "human-on-the-loop."

Affected rights: GDPR Art. 22 (right not to be subject to solely automated decision-making producing legal effects, without meaningful human involvement); right to an explanation and to contest a decision (GDPR Art. 13-15, 22(3)); DIFC Regulation 10.4 (named accountable person) effectiveness.

Evidence basis: Gap — "Human oversight degraded to rubber-stamp in practice" | status: in_progress | basis: EU AI Act Art. 14, 26 | inherent risk: yellow.

3.3 Risk: Incomplete special-category data safeguards

Description: It is unconfirmed whether CreditLens processes special-category data (e.g., health data referenced in lending/insurance contexts) or proxy variables correlating with protected characteristics, and safeguards for such data are only partially implemented.

Likelihood: Medium — the assessment session treats this conservatively as a partial (not confirmed absent) gap, pending a data-mapping exercise; the scope of exposure is genuinely unknown.

Severity: High — should special-category or proxy data be confirmed in scope for a credit-scoring model without adequate safeguards, this would implicate GDPR Art. 9 and heightened discrimination risk in a legal-effect decision context.

Affected rights: GDPR Art. 9 (special categories of personal data); UAE PDPL Art. 1, 10 (sensitive personal data); non-discrimination (DIFC Regulation 10.3).

Evidence basis: Gap — "Special-category data safeguards partial" | status: in_progress | basis: GDPR Art. 9; UAE PDPL Art. 1, 10 | inherent risk: yellow.

3.4 Risk: Insufficient data minimisation and retention controls

Description: Data minimisation and retention limits over inputs, prompts, and logs used by CreditLens are only partially configured, creating a risk of over-retention and processing beyond what is necessary for the credit-scoring purpose.

Likelihood: Medium — a partial (not absent) control exists per the assessment session, but no documented retention schedule is evidenced in the inputs.

Severity: Medium — over-retention of financial and personal data increases exposure in the event of a breach and is a distinct compliance gap from the DPIA and oversight risks above, though it does not itself produce an adverse decision.

Affected rights: GDPR Art. 5(1)(c) and (e) (data minimisation, storage limitation); UAE PDPL Art. 5 (processing principles).

Evidence basis: Gap — "Data minimisation/retention only partially configured" | status: open | basis: GDPR Art. 5; UAE PDPL Art. 5 | inherent risk: yellow.

3.5 Risk: Discriminatory or biased outcomes from inadequately governed training/input data

Description: Training and input data governance — relevance, representativeness, and provenance — is only partially documented, and bias testing has been performed once with no recurring schedule, creating risk that CreditLens produces discriminatory scoring outcomes against retail or SME clients, including on the basis of protected or proxy characteristics.

Likelihood: Medium — a single point-in-time bias test has occurred, meaning some assurance exists, but there is no ongoing monitoring to detect drift or emergent bias in a production credit-scoring context.

Severity: High — in a credit-scoring context, discriminatory outcomes directly affect access to essential financial services (declined facilities, reduced limits) and engage core non-discrimination protections.

Affected rights: Non-discrimination (DIFC Regulation 10.3; GDPR recital-level fairness principle, Art. 5(1)(a)); UAE PDPL processing-principles requirements; EU AI Act Art. 10 (data and data governance for high-risk systems).

Evidence basis: Gaps — "Bias testing done once, no recurring schedule" (basis: EU AI Act Art. 10; DIFC Reg. 10.3, status in_progress) and "Training/input data governance only partially documented" (basis: EU AI Act Art. 10, status in_progress).

3.6 Risk: Incomplete technical documentation and audit trail undermining accountability and data-subject redress

Description: Technical documentation for CreditLens is incomplete/outdated and the audit trail for decisions and model changes is only partially reconstructable, limiting the controller's ability to demonstrate compliance or respond to a data subject's request to understand or contest a decision.

Likelihood: Medium — these are confirmed, current gaps (not hypothetical), but do not by themselves produce an adverse decision; they impair the controller's ability to respond when one occurs.

Severity: Medium — where a data subject exercises a right to explanation or challenge under GDPR Art. 22(3) or equivalent, incomplete documentation and audit trail would impair Meridian's ability to provide a substantive response.

Affected rights: GDPR Art. 13-15 (information rights), Art. 22(3) (right to obtain human intervention, express point of view, contest decision); accountability principle (GDPR Art. 5(2)).

Evidence basis: Gaps — "Technical documentation incomplete/outdated" (basis: EU AI Act Art. 11, Annex IV,

status open) and "Audit trail only partially reconstructable" (basis: ISO/IEC 42001 Clause 8, status in_progress).

§ 4. Mitigation measures already in place

Based solely on gaps recorded as closed or on measures explicitly described as existing/operating in the assessment session:

- Documented lawful basis for processing is in place (DPS_01_YES), covering the underlying legal ground for CreditLens's processing of personal and financial data.
- Mapped and safeguarded cross-border transfers of this system's data are documented (DPS_04_YES), addressing DIFC-to-EU-to-UAE data flows implied by the Frankfurt branch structure.
- Human oversight structure is formally designed and operating (EUS_03_YES), with a named accountable person at the DIFC entity level (DIFS_03_YES), even though its practical effectiveness is separately identified as a residual risk (§ 3.2).
- DIFC autonomous systems register entry has been completed, satisfying DIFC Regulation 10.2 (DIFS_01_YES).
- Named, empowered system owner exists for governance purposes (GOV_01_YES).
- Production monitoring with defined metrics and alerts is operating (GOV_08_YES), providing visibility into performance degradation.
- Customer transparency channel: affected customers are informed that AI is involved and have access to a working explanation channel (GOV_07_YES), addressing baseline transparency expectations.
- One-time bias test has been performed on CreditLens (GOV_05_PART — partial, not a recurring programme, but an initial test has occurred).
- Kill-switch capability exists for the system, providing a baseline (though not fully developed) incident-response capability (GOV_06_PART).

No further mitigations are documented in the source assessment beyond those listed above.

§ 5. Residual risks

The following open or in-progress gaps from the assessment session remain unresolved and correspond directly to the risks identified in § 3:

Gap (ID reference by finding) · Status · Corresponding § 3 risk
No DPIA completed despite likely mandatory trigger (GDPR Art. 35; UAE PDPL Art. 21) · open · § 3.1
Human oversight degraded to rubber-stamp in practice (EU AI Act Art. 14, 26) · in_progress · § 3.2
Special-category data safeguards partial (GDPR Art. 9; UAE PDPL Art. 1, 10) · in_progress · § 3.3
Data minimisation/retention only partially configured (GDPR Art. 5; UAE PDPL Art. 5) · open · § 3.4
Bias testing done once, no recurring schedule (EU AI Act Art. 10; DIFC Reg. 10.3) · in_progress · § 3.5
Training/input data governance only partially documented (EU AI Act Art. 10) · in_progress · § 3.5
Technical documentation incomplete/outdated (EU AI Act Art. 11, Annex IV) · open · § 3.6
Audit trail only partially reconstructable (ISO/IEC 42001 Clause 8) · in_progress · § 3.6
AI-literacy training limited to general awareness (EU AI Act Art. 4) · open · § 3.2 (contributing factor to oversight degradation)
DIFC fairness/non-discrimination testing not systematic (DIFC Reg. 10.3) · open · § 3.5

None of the identified gaps in the source assessment are rated RED; all are rated YELLOW (inherent risk). No gap has yet been closed for these ten items — they remain open or in_progress as of the assessment completion date (2026-07-08). Until the DPIA (§ 3.1) is completed and the human-oversight redesign (§ 3.2) is implemented, the overall residual risk to data subjects from adverse, insufficiently-reviewed credit decisions should be treated as not yet reduced to an acceptable level, notwithstanding the mitigations in § 4.

§ 6. Consultation

6.1 DPO consultation

[To be completed by the controller] — the inputs do not confirm whether a Data Protection Officer has been appointed. This is explicitly listed as a "known unknown" in the profile summary ("Appointed governance roles: whether a Data Protection Officer (DPO) ... exists"). GDPR Art. 35(2) and DIFC Data Protection Law Art. 20(3) both require that, where a DPO is designated, their advice be sought in carrying out this DPIA. Meridian must confirm DPO status and, if appointed, record the DPO's input and sign-off on this document before finalisation.

6.2 Data subjects consultation

No evidence in the inputs indicates that data subjects or their representatives have been consulted on the CreditLens processing. DIFC Data Protection Law Art. 20(8) provides that a controller shall seek the input of data subjects or their representatives on intended processing "where appropriate," taking into account protection of commercial/public interests and processing security. [To be completed by the controller] — a determination of whether such consultation is appropriate for CreditLens and, if so, a record of the consultation undertaken.

6.3 Joint controllers

No joint-controller arrangement is described in the inputs. Meridian is identified as sole Provider/developer of CreditLens (AI builder archetype). [To be completed by the controller] — confirmation of whether any third party (e.g., bureau data providers, model-component vendors referenced in the CBUAE outsourcing-controls gap) acts as a joint controller or independent controller for any part of the data flows feeding CreditLens.

§ 7. Conclusion and sign-off

CreditLens is a high-risk, production, human-on-the-loop credit and counterparty scoring system that produces outputs with legal or similarly significant effects on retail and SME clients, including EU-domiciled clients served via the Frankfurt branch. This DPIA has been prepared to close the identified gap that "no DPIA has been conducted for a system that scores individuals and produces legal/significant effects" despite this being a likely mandatory trigger under GDPR Art. 35(3)(a), UAE PDPL Art. 21(2)(a), and DIFC Data Protection Law Art. 20.

The assessment identifies genuine compliance strengths — documented lawful basis, mapped cross-border transfer safeguards, a formally designed human-oversight structure, DIFC system registration, and production monitoring. However, six risks to data-subject rights remain open or in-progress, most critically the degradation of human oversight to a rubber-stamp function (§ 3.2) and the absence, until now, of a completed prior impact assessment (§ 3.1). Given the "high" user-impact level and Annex III high-risk classification of CreditLens, and the presence of unresolved YELLOW-rated gaps affecting automated decision-making safeguards and anti-discrimination controls, this processing should be considered to carry residual risk that has not yet been fully mitigated.

Per GDPR Art. 36 and the equivalent prior-consultation provisions referenced in the regulatory anchors, Meridian should assess whether prior consultation with the relevant supervisory authority is required, given that the DPIA — even after mitigation measures in § 4 are accounted for — does not yet demonstrate that all identified high risks (§ 3.1, § 3.2, § 3.5) have been reduced to an acceptable level.

Sign-off:

- Controller representative: [To be completed by the controller]
- DPO (if appointed): [To be completed by the controller]
- Date: [To be completed by the controller]

§ 8. Regulatory anchors (verbatim)

GDPR Art. 35

"1. Where a type of processing in particular using new technologies, and taking into account the nature, scope, context and purposes of the processing, is likely to result in a high risk to the rights and freedoms of natural persons, the controller shall, prior to the processing, carry out an assessment of the impact of the envisaged processing operations on the protection of personal data. A single assessment may address a set of similar processing operations that present similar high risks.

2. The controller shall seek the advice of the data protection officer, where designated, when carrying out a data protection impact assessment.

3. A data protection impact assessment referred to in paragraph 1 shall in particular be required in the case of:

(a) a systematic and extensive evaluation of personal aspects relating to natural persons which is based on automated processing, including profiling, and on which decisions are based that produce legal"

[Source: EU-GDPR: Regulation on general data protection — chunk 118/119]

GDPR Art. 36 (Prior consultation)

"1. The controller shall consult the supervisory authority prior to processing where a data protection impact assessment under Article 35 indicates that the processing would result in a high risk in the absence of measures taken by the controller to mitigate the risk.

2. Where the supervisory authority is of the opinion that the intended processing referred to in paragraph 1 would infringe this Regulation, in particular where the controller has insufficiently identified or mitigated the risk, the supervisory authority shall, within period of up to eight weeks of receipt of the request for consultation, provide written advice"

[Source: EU-GDPR: Regulation on general data protection — chunk 126]

DIFC Data Protection Law Art. 20 (referenced as "Art. 40" context / DIFC Regulation 10.34 equivalent)

"the risks to the rights of the Data Subjects concerned. A Controller may also elect to carry out such assessment in relation to the Processing of Personal Data that is not a High Risk Processing Activity.

(2) A single assessment may address a set of similar Processing operations that present similar risks... (3) A DPO, where appointed, shall be responsible for overseeing data protection impact assessments. (4) The Commissioner may at his discretion publish a non-exhaustive list of types or categories of Processing operations that are considered to be High Risk Processing Activities. Such a list is not"

[Source: Data Protection Law (DIFC) — chunk 55]

"Processing operations that are considered to be High Risk Processing Activities. Such a list is not intended to

be exhaustive and does not absolve a Controller from responsibility for complying with this Law in all respects with regard to High Risk Processing Activities.

(5) The Commissioner may also publish a list of the types or categories of Processing operations for which no data protection impact assessment is required.

(6) A data protection impact assessment shall contain at least:

(a) a systematic description of the foreseen Processing operations and the purpose(s) of the Processing, including, where applicable, the legitimate interest pursued by a Controller;

(b) an assessment of the necessity and proportionality of the Processing operations in relation to the purpose(s);

(c) identification and consideration of the lawful basis for the Processing, including:"

[Source: Data Protection Law (DIFC) — chunk 56]

"this Law, taking into account the rights and legitimate interests of Data Subjects and other concerned persons.

(7) In assessing the impact of the Processing operations, compliance with approved codes of conduct referred to in Article 48 by a Controller or Processor shall be taken into account.

(8) Taking into account protection of commercial or public interests or the security of Processing operations, a Controller shall seek the input of Data Subjects or their representatives on the intended Processing, where appropriate."

[Source: Data Protection Law (DIFC) — chunk 58]

DIFC Regulation 10 (ADGM Data Protection Regulations 2021, cross-referenced)

"34. Data Protection Impact Assessment

(1) The Controller must, prior to Processing that is likely to result in a high risk to the rights of natural persons, carry out an assessment of the impact of the envisaged Processing operations on the protection of Personal Data (a 'Data Protection Impact Assessment').

(2) A single Data Protection Impact Assessment may address a set of similar Processing operations that present similar high risks. The outcome of the Data Protection Impact Assessment should be taken into account when determining the appropriate measures to be taken in order to demonstrate that the Processing of Personal Data complies with these Regulations.

(3) The Controller must seek the advice of the Data Protection Officer, where designated, when carrying out a Data Protection Impact Assessment."

[Source: ADGM Data Protection Regulations 2021 — chunk 108]

"(5) The Data Protection Impact Assessment must:

(a) describe the nature, scope, context and purpose of the Processing;

(b) assess necessity, proportionality and compliance measures;

(c) identify and assess risks to individuals; and"

[Source: ADGM Data Protection Regulations 2021 — chunk 109]

UAE PDPL Art. 21

"Article 21

Assessment of the Impact of Personal Data Protection

1. Taking into account the nature, scope and purposes of data processing, the Controller shall, before carrying out the processing, evaluate the impact of the proposed processing operations on the protection of Personal Data, when using any of the modern technologies that would pose a high risk to the privacy and confidentiality of the Data Subject's Personal Data.

2. The assessment of the impact provided for in Paragraph (1) of this Article shall be required in the following cases:"

[Source: UAE Federal PDPL (Decree-Law 45/2021) — chunk 46]

"a. If the processing includes a systematic and comprehensive assessment of the personal aspects of the Data Subject, using automated processing, including profiling, having legal consequences or serious impact on the Data Subject.

b. If processing would be carried out on a large volume of Sensitive Personal Data.

3. The assessment stipulated in Paragraph (1) of this Article shall include, at a minimum, the following:

a. Clear and systematic explanation of the suggested processing operations for the protection of Personal Data and the purpose of processing.

b. Evaluation of how necessary the processing operations are and how they are suitable for the purpose of processing.

c. Evaluation of potential risks related to the privacy and confidentiality of the Data Subject's Personal Data.

d. The suggested procedures and measures aimed at reducing the potential risks related to"

[Source: UAE Federal PDPL (Decree-Law 45/2021) — chunk 47]

§ 9. Document metadata

- DPIA subject system: CreditLens (v4.1), system ID f4fb1f74-3f52-4298-911b-1a8c78f3e513
- Controller: Meridian Capital Partners
- Assessment session referenced: 123fe04f-0676-401b-bed5-4b38a7256a6f, completed

2026-07-08T17:26:38.224+00:00, assessor: GUARD GRC

- Regulatory regimes in scope for this DPIA: GDPR, DIFC Regulation 10 / DIFC Data Protection Law, UAE Federal PDPL
- Systems explicitly out of scope for this DPIA: Advisory chatbot (retail investment guidance), third-party CV-screening tool (HR recruitment) — both flagged in the source assessment as requiring separate assessment. Shadow-AI use of general-purpose AI assistants by employees is also out of scope for this document and is noted only as an organizational governance gap disclosed in the profile summary.
- DPIA author: [To be completed by the controller] — name/role of DPIA preparer within Meridian or its advisers
- Review trigger: Per GDPR Art. 35 and DIFC Data Protection Law Art. 20(10), this DPIA should be reviewed on a regular basis proportionate to the extent and type of processing, or when there is a change in the risk represented by the processing operations (e.g., model retraining, new data sources, expansion to new jurisdictions).
- Version: 1.0 (initial DPIA, prepared to close open gap "No DPIA completed despite likely mandatory trigger")
- Distribution: [To be completed by the controller] — internal governance committee, DPO (if appointed), DFSA/CBUAE as required by prior-consultation analysis in § 7

Integrity

SHA-256 hashes of every document embedded in this pack, computed over the stored content at generation time:

Policy · AI Governance and Human Oversight Policy
9335d5c63ee9d0d270ca0e27e1af242ce61bc26bb9528d664dfb22efea0784c0

Policy · AI Risk, Incident Response, and Data Breach Management Policy
9b1acce974b955c45dba05b94cf593971fbcfc3d5ab09a499d0ce04e12ea1765

Assessment report · Meridian Capital Partners
6a17f7a46c1a04bc2d95f843e6f30b4e5872f7ae090ce5554d2559d3976f2af2

Assessment report · Meridian Capital Partners
c7fa488b634cc1bffe4a7fb026b29e2c9c4f230c832352bc92b83037cf896699

AI Transparency Statement · AI Transparency Statement — Meridian Capital Partners
192b3821e72218f6d102e06cd6b2c32bf972718faf8a6520365ebffcc93dbf15

Privacy Notice · Privacy Notice — Meridian Capital Partners
1bafc4a2de9762cf9defab2198f5e7e3b03dad8a177eace87ebb306472475340

Data Protection Impact Assessment · Data Protection Impact Assessment: CreditLens Credit and Counterparty Scoring Model
d3f5c3afd93028fb93089cec5f419e01307bc6362b8e46e4c1088d4b73d4c657

Pack PDF SHA-256: recorded in the GUARD governance register at generation time and shown alongside every download and share view.

This pack is a point-in-time compilation of governance records held in GUARD GRC at the generation timestamp, assembled for regulatory correspondence. Embedded documents are first-draft governance artefacts reviewed and approved through the organisation's sign-off workflow where so stated. GUARD GRC L.L.C-FZ provides no legal warranty.