

Project Meridian

AI & Technical Due Diligence — Red-Flag Screen

YSA Systems Inc.

Yuri Subach (yuri@ysasystems.com)

June 25, 2026

1 Engagement details

- *Target:* Project Meridian — [AI document-processing platform: intelligent extraction from enterprise business documents]
- *Prepared for:* [Private-equity sponsor] — Investment Team
- *Engagement type:* Red-Flag Screen
- *Engagement reference:* R26-A111-Sample
- *Classification:* Confidential — Privileged
- *Notice:* This report is confidential and privileged. It is prepared solely for the named recipient and may not be reproduced, distributed, or relied upon by any other party without prior written consent. Findings reflect information available as of the date above and the agreed scope of work.

Illustrative sample: Client, target, personnel, and identifying details have been redacted or fictionalised.

2 Executive summary

Overall technical risk: Elevated. One or more material concerns are present that full assessment should verify and quantify. Risk is concentrated in the AI/ML domain; the surrounding application and infrastructure do not raise material concern at this stage.

Overall: Elevated

I



This screen was engaged to identify, ahead of a Letter of Intent, any technical matter material enough to warrant the sponsor’s attention before committing to full assessment on Project Meridian — an AI document-processing platform that extracts structured data from unstructured business documents. Work was time-boxed to five business days against a limited materials set and management access.

The target’s conventional software engineering is sound. The concerns concentrate in the asset the thesis rests on — the extraction capability itself. On the evidence available, that capability is *substantially rented rather than owned*, its training-data provenance is undocumented, and the pipeline that produces it cannot presently be reproduced by the current team. These are not defects that block a transaction, but each bears directly on how defensible the core asset is, and each is material enough that the sponsor will want it resolved before relying on the AI as a durable advantage.

2.1 Most material findings

Top findings are summarized below; the full set is in Section 5.

- **Critical (AI-01)** — Core inference depends on a single third-party model provider, with no abstraction layer or tested fallback. Capability, unit economics, and continuity are exposed to one vendor’s pricing and terms.
- **Critical (AI-02)** — Training and fine-tuning data provenance is undocumented; part of the corpus appears to have been assembled without a clear licensing or consent trail, and may contain personal data.
- **Material (AI-03)** — The fine-tuning pipeline was built and operated by a since-departed engineer; the current team could not demonstrate an end-to-end retrain. The model is not presently reproducible.
- **Material (AI-05)** — Marketed extraction accuracy (“~99%”) is not supported by any independent held-out evaluation; observed performance varies materially across document types.
- **Material (SC-01)** — The SOC 2 referenced in sales materials is a lapsed Type I, not a current Type II. Regulated enterprise customers are likely to require the latter.

Where the risk concentrates. The central open question is the *defensibility of the AI asset*: how much of the extraction capability the company actually owns, where its training data came from, and whether it can be rebuilt without the original author. Full assessment should be scoped around resolving those three questions first (see Section 6). Findings are stated as technical facts; decisions on how they bear on valuation, terms, or whether to proceed rest with the sponsor and its advisors.


3 Risk summary

3.1 Risk overview

Findings by severity:



Number of findings per review area:



3.2 Findings by domain

ID	Domain	Finding	Severity
AI-01	Third-party AI dependence	Single-provider inference dependency, no fallback	Critical
AI-02	Data provenance & rights	Undocumented training-data lineage; licensing & PII exposure	Critical

ID	Domain	Finding	Severity
AI-03	Model ownership & reproducibility	Pipeline not reproducible; key-person dependency	Material
AI-04	MLOps maturity	No evaluation harness or drift monitoring; manual deploys	Material
AI-05	Capability credibility	Marketed extraction accuracy unsupported by held-out evidence	Material
SC-01	Security & compliance	SOC 2 posture overstated; control gaps observed	Material
IP-01	IP & open-source	Copyleft component in the serving path	Material
AI-06	Model generalization	Extraction tuned narrowly to a few document layouts	Material
PA-01	Product & architecture	Conventional, sound application architecture	Noted

3.3 Definitions

Overall technical-risk scale:

- **Low** — no material technical concerns in the areas screened.
- **Moderate** — some concerns, none individually material on the evidence seen.
- **Elevated** — one or more material concerns full assessment should verify and quantify. (*This engagement*)
- **High** — concern(s) severe enough to be central to the client’s own evaluation of the asset.

Finding severity (RAG):

- **Critical** — potential to materially affect the technical soundness or value of the asset; described so the client and its advisors can weigh it.
- **Material** — a real concern requiring full assessment; likely manageable but must be scoped.
- **Noted** — informational; no material technical risk identified.

4 Scope, basis & limitations

4.1 Objectives

1. Is the extraction capability genuinely the company’s own, and how defensible is it?

2. Where does the training data come from, and does the company have the right to use it?
3. Does the product perform as represented?
4. Are there technical, security, or compliance conditions likely to matter to the target's customers or to integration?

4.2 Materials reviewed

1. Data room: architecture overview, security questionnaire, third-party vendor list, one SOC 2 report, sales collateral.
2. Two management calls — [CTO] and acting head of engineering (around 90 minutes each).
3. Live product demonstration on target-supplied sample documents.
4. Read-only walkthrough of selected repositories and infrastructure-as-code.

4.3 Effort and window

- Five business days. This is a screen, not an audit: depth was traded for speed by design, to inform the LOI decision economically.

4.4 Explicitly out of scope

- Full source-code audit and independent security testing (no penetration test).
- Independent model evaluation on a controlled held-out set.
- Financial, commercial, and legal diligence.

4.5 Basis and limitations

- Findings rest on materials provided and management representations, not independently audited except where stated.
- Access was constrained to what is appropriate at screen stage: no full code review, no code execution, no production or customer-data access.
- Point-in-time assessment; systems and third-party services may change after the report date.
- A technical assessment only. Not an audit, and not legal, tax, or financial advice; not a recommendation on whether, at what value, or on what terms to transact.
- Prepared solely for the named recipient; not to be relied upon by any other party.

5 Detailed findings

5.1 Core extraction depends on a single external model provider

ID: AI-01 Domain: Third-party AI dependence Severity: Critical

Observation The document-understanding capability the product is sold on is delivered by prompting a single commercial large-language-model API. There is no provider-abstraction layer, no tested secondary provider, and no self-hosted fallback. The company’s own contribution sits in prompt orchestration, post-processing, and the application around the model — not in the model itself.

Basis Vendor list and architecture overview (data room); confirmed on the CTO call, where the provider was named as the sole inference dependency and no fallback was reported as built or tested.

Impact Three coupled exposures: *economic* — gross margin is a function of a price the company does not control; *continuity* — a provider outage, rate-limit change, or policy change directly interrupts the product; and *defensibility* — a capability that is rented is available to competitors on the same terms. In substance the company is closer to an application layer over a shared model than an owner of differentiated AI.

Follow-up Confirm contract terms, pricing tier, and rate limits with the provider; quantify inference cost as a share of gross margin; test portability to an alternative or self-hosted model, and at what accuracy and cost.

5.2 Training-data provenance is undocumented

ID: AI-02 Domain: Data provenance & rights Severity: Critical

Observation The fine-tuning corpus used to adapt the model to the target’s document types has no data-lineage record. Management could not produce a documented chain of custody, licensing terms, or consent basis. On the call, part of the corpus was described as customer documents “collected over time” and retained beyond the processing purpose for which they were submitted — business documents that routinely contain personal data.

Basis CTO and head-of-engineering calls; absence of any data-provenance or data-processing documentation in the data room.

Impact Unclear rights in training data is a defensibility and exposure question at once: it affects whether model outputs are cleanly ownable, whether the company can represent clear title to the asset, and whether reuse of customer documents for training is consistent with its own customer agreements and applicable privacy law. Provenance gaps are difficult to remediate retroactively once a model has been trained on the data.

Follow-up Reconstruct data lineage per corpus component; establish the licensing/consent basis for each source; identify personal data and its handling; where lineage cannot be established, scope what retraining on clean, documented data would require.

5.3 The model pipeline is not presently reproducible

ID: AI-03 Domain: Model ownership & reproducibility Severity: Material

Observation The fine-tuning and evaluation pipeline was built and operated by a lead ML engineer who has since left the company. It is not captured as versioned, runnable code; training configuration and data-preparation steps live partly in that individual's notebooks and partly in undocumented convention. The current team could not demonstrate an end-to-end retrain during the engagement.

Basis Head-of-engineering call; repository walkthrough showed notebook-based, un-versioned training artifacts rather than a reproducible pipeline.

Impact If the model cannot be retrained, the company cannot reliably improve it, recover it if a hosted artifact is lost, or migrate it off the current provider (compounding AI-01). It also concentrates critical knowledge in a departed individual — a key-person dependency that has already partly materialised.

Follow-up Require a full retrain from documented inputs; assess the effort to rebuild the pipeline as versioned, reproducible code; confirm the departed engineer's access and retained knowledge.

5.4 No evaluation harness or production model monitoring

ID: AI-04 Domain: MLOps maturity Severity: Material

Observation There is no automated evaluation harness scoring model changes against a fixed benchmark before release, no regression testing on prompt or model updates, and no production monitoring of extraction quality or drift. Changes are applied manually and validated by spot-checking.

Basis CTO call; no evaluation or monitoring tooling evident in the repositories or infrastructure reviewed.

Impact Without an eval harness the company cannot tell whether a change improves or degrades quality, and cannot safely respond to provider-side model updates, which can shift extraction behaviour silently. Quality regressions can reach customers undetected — consequential where extracted data feeds downstream business processes.

Follow-up Assess the effort to stand up an eval harness and drift monitoring; review incident history tied to undetected quality regressions.

5.5 Marketed accuracy is not supported by independent evidence

ID: AI-05 Domain: Capability credibility Severity: Material

Observation Sales materials cite extraction accuracy of approximately 99%. The only supporting evidence is a company-run benchmark on a favourable, self-selected sample. In the live demo on the target’s own documents, performance was serviceable but visibly lower and varied by document type — weaker on lower-quality scans and less common layouts. No independent held-out evaluation exists, and no field-level accuracy breakdown was available.

Basis Sales collateral (data room); direct observation during the product demonstration.

Impact The gap between the marketed figure and observed performance affects how the capability should be understood and where it is reliable. A single headline accuracy number is of limited meaning in extraction, where per-field accuracy and error mode matter more than an aggregate; the figure is best treated as unverified until measured independently.

Follow-up Run an independent evaluation on a controlled, representative held-out set spanning document types and quality levels; report per-field accuracy with variance rather than a single aggregate.

5.6 Compliance posture is overstated relative to evidence

ID: SC-01 **Domain:** Security & compliance **Severity:** Material

Observation Sales materials reference “SOC 2” without qualification. The report in the data room is a lapsed Type I from an earlier period — a point-in-time design assessment, not the Type II operating-effectiveness attestation enterprise buyers typically require. The questionnaire and call also surfaced control gaps: secrets handling relies partly on developer convention, and access reviews are ad hoc rather than on a defined cadence.

Basis SOC 2 report and security questionnaire (data room); CTO call.

Impact The gap between represented and actual compliance status is itself a finding. Separately, a current Type II is often a precondition for enterprise customers entrusting sensitive documents to a third-party processor, so the gap may bear on the target’s ability to move up-market. The observed control gaps are common at this stage and remediable.

Follow-up Confirm current attestation status and the path and effort to Type II; validate the specific control gaps and remediation cost. Contractual representations about certifications are for the sponsor’s counsel.

5.7 A copyleft component sits in the serving path

ID: IP-01 **Domain:** IP & open-source licensing **Severity:** Material

Observation Dependency review surfaced a component under a strong copyleft (AGPL-class) license in the document ingestion and inference path. Whether its obligations are triggered depends on the specific license version and how the component is deployed and offered to users — a fact-specific question the company had not assessed.

Basis Repository walkthrough and dependency manifest review.

Impact If obligations apply, they could be inconsistent with a wholly-proprietary IP position for the affected part of the stack. Flagged as a technical fact for legal assessment; typically resolvable by replacing the component or satisfying the license terms.

Follow-up Produce a full open-source bill of materials with licenses; isolate the component and deployment mode for counsel; scope replacement effort if needed.

5.8 Extraction is tuned narrowly to a few document layouts

ID: AI-06 **Domain:** Model generalization **Severity:** Material

Observation The model has been tuned around the document templates of a small number of early customers, whose material dominates the fine-tuning data. Performance on those layouts is the basis of the accuracy claims; performance on materially different layouts — as would arrive with new customers — is untested.

Basis Head-of-engineering call on fine-tuning data composition; observed variance across document types in the demo (relates to AI-05).

Impact Purely technical: capability demonstrated on current customers may not transfer to new ones without further tuning, which affects the engineering effort behind onboarding and expansion. Stated as a generalization property of the model, not as a commercial judgment about the customer base.

Follow-up Evaluate on document layouts outside the current customer set; estimate the per-customer tuning effort new formats would require.

5.9 Application architecture is conventional and sound

ID: PA-01 **Domain:** Product & architecture **Severity:** Noted

Observation Setting the model aside, the surrounding product — ingestion, queueing, storage, review UI, and API — is built on a standard, well-understood cloud-native stack (Python server-side, PostgreSQL, React front-end) with sensible service boundaries and partial infrastructure-as-code. Nothing in the architecture review raised a material concern at screen level.

Basis Architecture overview (data room); repository and IaC walkthrough.

Impact None material at this stage.

Follow-up Standard code-quality and scalability review.

6 Suggested full assessment scope

If the sponsor advances, the following technical matters — ordered by materiality — are what a full assessment should verify and quantify. Recommending the technical work is within remit; whether to commission it, and how these findings bear on the transaction, is the sponsor’s decision.

1. **Establish what the company actually owns.** Quantify inference cost as a share of margin; test portability to an alternative or self-hosted model. *Addresses AI-01.*
2. **Reconstruct training-data provenance.** Build the data-lineage record, establish licensing/consent basis per source, identify personal data; scope clean-retrain effort where lineage fails. *Addresses AI-02; rights questions to counsel.*
3. **Prove reproducibility.** Require a full end-to-end retrain from documented inputs; assess effort to rebuild the pipeline as versioned code. *Addresses AI-03.*
4. **Measure the capability independently.** Evaluate on a controlled held-out set across document types and quality levels, reporting per-field accuracy and variance. *Addresses AI-05 and AI-06.*
5. **Verify compliance status and controls.** Confirm current attestation and path to Type II; validate control gaps and remediation cost. *Addresses SC-01.*
6. **Complete the open-source bill of materials.** Isolate the copyleft component and deployment mode for counsel; scope replacement. *Addresses IP-01.*
7. **Stand up MLOps basics.** Scope an evaluation harness and drift monitoring; review regression-related incident history. *Addresses AI-04.*

Items 1–3 address the defensibility of the core AI asset and are the highest priority; the balance can proceed in parallel.

— *End of report* —