

ILLUSTRATIVE SAMPLE / THE STANDARD V3.1

# AI Trust File

*An evidence-based record of how an organization governs its artificial intelligence.*

This file records what was examined, what the evidence shows, and the organization's result against The Standard, reviewed and scored on a 100-point scale across four pillars of responsible AI.

PREPARED IN CONFIDENCE FOR THE BOARD  
OF

**Avante**

FILE  
QIA-2026-007

ISSUED  
9 June 2026

# Contents

---

01	Executive summary	05
02	Why this result	06
03	Recommended board actions	07
04	Scope & methodology	08
05	AI inventory & use-case map	09
06	Score breakdown by pillar	10
07	Evidence-backed findings	11
08	Material issue, gaps & remediation	12
09	Source coverage	13
10	Evidence request & next level	14
11	Regulatory alignment map	15
12	Evidence Register	16
13	Verification & disclosure	17
14	Attestation & confidence	18
15	Suspension & disclaimer	19

THE RESULT

---

**82** / 100

---

REVIEW TIER COMPLETED

**Gold**

REVIEWER  
CONFIDENCE

**High**

OPPORTUNITIES  
NOTED

**3**

Evidence basis E2, sample

# Illustrative sample

Diagnostic completed

RECOMMENDED RELIANCE

**Illustrative sample, not for reliance**

REVIEW TIER  
COMPLETED

**Gold**

NEXT REVIEW

**Sample**

This is an illustrative sample built to show the AI Trust File format. Avante is a fictional organization. Within the stated scope, Avante presents a documented, human-led AI-governance posture with credible controls over its credit-decisioning systems. The open items are quantitative bias evidence and broader external assurance, not missing fundamentals.

TRUST STRENGTHS

A standing AI governance committee drawn from Risk, Legal, and Data, with a named owner per system.

A documented, mandatory human review before any adverse credit decision.

Encryption and role-based access controls evidenced across the AI life cycle.

POINTS TO RESOLVE

Quantitative fairness testing is documented in method but not yet published in results.

Customer-facing explainability is available on request rather than by default.

External AI-specific assurance is in progress rather than complete.

REQUIRED ACTIONS

Publish summary results of quantitative bias testing for deployed credit models.

Make decision-factor explanations available to customers by default.

Complete external ISO/IEC 42001 alignment for the AI management system.

CONDITIONS BEFORE RELIANCE

This is an illustrative sample for a fictional organization; it is not issued and is not for reliance.

EVIDENCE STILL NEEDED

Published bias-testing results, the full model inventory, and the completed external alignment report.

# Executive summary

---

Review status **illustrative sample** · fictional organization

**G**rupo Financiero Avante governs the AI inside its consumer-credit and customer-service operations with a documented, human-led program. Within the stated scope it presents a well-structured posture; the open items are quantitative fairness evidence and broader external assurance.

The review examined Avante's AI use policy and governance framework, the AI systems declared in its credit and service operations, and the controls evidenced over data and decisions. It found a documented, human-led program with credible technical controls and a clear path to broader external assurance.

*“A documented, human-led AI program with credible controls over consequential decisions.”*

From here, publishing fairness results and completing external alignment would consolidate an already strong position.

# Why this result

---

Avante scored 82 because its AI governance is documented and human-led and its core controls are evidenced. Pending publication of fairness results and completion of external assurance caps transparency, not the fundamentals.

## WHAT EARNED THE RESULT

A documented AI use policy and cross-functional oversight committee.

A mandatory human-in-the-loop control for adverse decisions.

Evidenced encryption and access controls across the data flow.

## OPPORTUNITIES TO REACH THE NEXT LEVEL

Publish quantitative bias-testing results for deployed models.

Provide customer-facing explanations by default.

Complete ISO/IEC 42001 external alignment.

## WHAT CAPPED THE RESULT

Transparency is the main limit: fairness methods are documented but results are not yet published, and default customer explanations are not yet in place. Publishing both would raise the result.

# Recommended board actions

---

The actions below are ordered by priority for the board.

**1** Approve publication of bias-testing summaries for deployed models.

---

**2** Adopt default customer-facing explanations for AI decisions.

---

**3** Complete external ISO/IEC 42001 alignment.

---

**4** Maintain the model inventory and annual control reviews.

---

# Scope & methodology

## SYSTEMS REVIEWED

SYSTEM	PURPOSE	REVIEWED
<b>Credit decisioning</b>	Automated scoring support for consumer-credit decisions	Reviewed, supported by public evidence
<b>Customer-service assistant</b>	AI support within customer-service channels	Reviewed, supported by public evidence

## HOW THE REVIEW IS RUN

The Quiet AI review is run end to end by an automated diagnostic engine. It performs the analysis, scores every criterion against the standard, and verifies its own work. Three controls govern it:

### Evidence enforcement

No score is awarded without a specific supporting source. Absent evidence is recorded as insufficient and scored conservatively.

### Three-pass verification

Each finding is drafted, challenged by a second pass, then confirmed by a third that checks every citation against its source.

### Conservative scoring

Where evidence is ambiguous or incomplete, the lower score is recorded. The engine is built to under-state rather than over-state.

## HOW TO READ THE SCORE

The review tier sets how deep the evidence goes; the score is a diagnostic reference only. Reviewer confidence is recorded separately from the score, so a strong score resting on limited evidence is shown as such.

# AI inventory & use-case map

The AI systems in scope, their risk profile, and the evidence behind each. For public-source reviews, unknown fields are shown plainly rather than assumed.

SYSTEM	PURPOSE & USERS	DATA	HUMAN OVERSIGHT	RISK CLASS	EVIDENCE
<b>Credit decisioning model</b>	Consumer-credit scoring support; external	Personal and financial data	In-the-loop; human decision required	High-impact	E2
<b>Customer-service assistant</b>	Customer interactions; external	Personal data	On-the-loop; human handoff	Limited	E2
<b>Internal analytics support</b>	Operational analytics; internal	Aggregated data	On-the-loop	Limited	E2

## RISK CLASSES & FLAGS

**Risk:** minimal · limited · transparency · high-impact · unacceptable candidate · unknown. **Flags:** GenAI · autonomous agent · RAG · security-critical · workforce · sensitive-document · model-provider dependency.

# Score breakdown by pillar

Four pillars, each worth 25 points, for a total of 100. Score bands are diagnostic references only; they do not determine whether a TrustFile Certificate is issued.

PILLAR	SCORE	OF PILLAR
<b>I Fairness &amp; Non-Discrimination</b> bias testing across protected characteristics; training-data representativeness; disparate-impact monitoring; mitigation	20 / 25	80%
<b>II Transparency &amp; Explainability</b> system documentation; decision explainability; automated-decision disclosure; logging & traceability	21 / 25	84%
<b>III Data Governance, Privacy &amp; Security</b> provenance & lineage; lawful basis; minimization & retention; security; breach handling	20 / 25	80%
<b>IV Accountability &amp; Oversight</b> ownership & roles; human oversight; reliability & safety (output-accuracy & content-safety); risk management; monitoring & redress	21 / 25	84%
<b>Overall</b>		<b>82 / 100</b>
Fairness & Non-Discrimination		20
Transparency & Explainability		21
Data Governance, Privacy & Security		20
Accountability & Oversight		21
SECTOR BENCHMARK (ILLUSTRATIVE)		
<b>Avante</b>		<b>82</b>
Sector leaders		78
Sector average		45

# Evidence-backed findings

---

One finding per criterion, each with a score, a verified excerpt from public sources, and its origin.

---

## I.1 Bias testing across protected characteristics

**STRONG · E2**

Avante's published AI Use Policy requires a documented human review before any adverse credit decision is communicated to a customer.

SOURCE [Avante AI Use Policy \(sample\)](#) · EVIDENCE E2

---

## IV.2 Human oversight of decisions

**PARTIAL · E2**

The governance framework names an accountable owner for each AI system and a standing review committee drawn from Risk, Legal, and Data.

SOURCE [Avante governance framework \(sample\)](#) · EVIDENCE E2

---

## III.4 Security controls (encryption & access)

**STRONG · E2**

Encryption in transit and at rest, and role-based access control, are evidenced across the credit-decisioning data flow.

SOURCE [Avante security controls \(sample\)](#) · EVIDENCE E2

---

EVIDENCE MATURITY E0 unsupported · E1 self-attested · E2 public · E3 client document · E4 tested or corroborated · E5 monitored

# Material issue, gaps & remediation

## MATERIAL ISSUE

No single material issue surfaced within scope. The opportunities below are improvements, not deficiencies.

## OPPORTUNITIES TO STRENGTHEN

### Quantitative fairness evidence PRIORITY

Publish summary results of bias testing for deployed credit models.

### Default explainability PRIORITY

Provide decision-factor explanations to customers by default rather than on request.

## REMEDIATION TIMELINE

WHEN	ACTION	EVIDENCE PRODUCED
0-30d	Publish a customer-facing summary of how AI supports credit decisions.	Customer AI notice
30-90d	Publish quantitative bias-testing results for deployed credit models.	Bias-testing summary
90d+	Complete external ISO/IEC 42001 alignment for the AI management system.	ISO 42001 alignment report

# Source coverage

What the review searched, and what it found. Making the evidence boundary explicit allows the result to be interpreted responsibly.

SOURCE CATEGORY	REVIEWED	RESULT	TYPE
Trust centre & security page	yes	found	Public
Compliance & certifications	yes	found	Public
Responsible-AI / AI policy	yes	found	Public
Privacy notice	yes	found	Public
Terms of service	yes	found	Public
Product & documentation	yes	found	Public
Hiring & organisational signals	yes	partial	Public
Public reviews & reputation	yes	found	Public
Regulatory / adverse search	yes	none found	Public
Third-party & subprocessors	yes	partial	Public

**Review status** : company participated **no** · authorized **no** · internal documents reviewed **no** · controls tested **no** · monitoring active **no**. Source code and proprietary internals are never required.

# Evidence request & path to the next level

The specific evidence that would raise the score or reviewer confidence, and the review tier at which each piece is incorporated.

EVIDENCE REQUESTED	WHY IT MATTERS	REVIEWED AT	AFFECTS	PRIORITY
<b>AI Use Policy (full text)</b>	Confirms the governance baseline · Pillar IV	Gold	Confidence	High
<b>AI system inventory</b>	Confirms what is in production · Pillar I	Gold	Confidence	High
<b>Bias-testing results (deployed models)</b>	Substantiates fairness claims · Pillar I	Platinum	Score + confidence	High
<b>Customer explainability artefacts</b>	Strengthens transparency · Pillar II	Platinum	Score + confidence	High

## THE PATH

### Silver

Online search + self-attestation. Baseline diagnostic.

### Gold

Silver + questionnaire. Public evidence plus client answers.

### Platinum

Gold + document upload. Evidence incorporated into the review.

### Diamond

Platinum plus monthly monitoring within defined scope.

# Regulatory alignment map

This review maps to the frameworks below. Alignment indicates mapped evidence within the review scope; it is not a certification of legal compliance.

FRAMEWORK	ALIGNMENT	NOTES
EU AI Act	PARTIAL	Credit decisioning is treated as high-impact under the EU AI Act, whose high-risk obligations phase in through 2026 to 2027; human review and risk controls are documented, external conformity work is in progress.
ISO/IEC 42001	ALIGNED	Governance mirrors an AI management system aligned to ISO/IEC 42001; external alignment is in progress.
NIST AI RMF	ALIGNED	Risk-assessment and oversight practices map to the NIST AI RMF functions of govern, map, measure, and manage.
GDPR	ALIGNED	A privacy program evidences lawful basis, minimization, and data-subject rights consistent with the GDPR.
PIPEDA	ALIGNED	Canadian operations evidence consent, accountability, and safeguards consistent with PIPEDA.
CCPA / CPRA	PARTIAL	Automated credit decisions are in scope of the CCPA as amended by the CPRA, whose ADMT and risk-assessment rules apply from 2026; default opt-out and disclosures are in progress.
LFPDPPP (2025)	ALIGNED	Privacy notices, ARCO rights, and controller duties evidence practice under Mexico's 2025 LFPDPPP.

# Evidence Register

---

Every score traces to a row here. For an illustrative sample, all sources shown are public.

ID	SOURCE	EVIDENCE (EXCERPT)	SUPPORTS	CONFIDENCE
E-1	Avante AI Use Policy (sample)	Human review required before adverse decisions.	III.1 Data provenance & lineage	High
E-2	Avante governance framework (sample)	Named owner and review committee per AI system.	IV.2 Human oversight	High
E-3	Avante security controls (sample)	Encryption in transit and at rest evidenced.	III.4 Security controls	Medium
E-4	Avante AI Use Policy (sample)	Permitted and prohibited AI uses defined.	II.2 Decision explainability	High

# Verification & disclosure

---

## PUBLIC VERIFICATION RECORD

**Registry status** Sample only, no certificate issued

---

**Valid through** Not applicable, illustration

---

**Verification URL** [quietaistandard.org/verify/avante.mx](https://quietaistandard.org/verify/avante.mx) (sample)

---

*A public verification record may be issued for any completed paid TrustFile Certificate if the client elects public disclosure.*

## DISCLOSURE STATUS

PUBLIC

CONFIDENTIAL

---

Certificate status, review tier, validity (score optional)

Intake questionnaire and submitted documents

---

Verification record (when disclosure elected)

Internal evidence not cleared for release

---

# Attestation & confidence

---

## CONFIDENCE BASIS

Confidence is High for this illustrative sample, reflecting a complete and internally consistent evidence set within the defined scope. In a live review, publishing fairness results and completing external alignment would further strengthen the basis of the result.

## REVIEW BASIS

This file was produced under one of the bases below; the applicable basis is recorded on the result page.

**Authorized paid review:** commissioned by the organization, which represented that intake information was accurate and that it was authorized to submit it.

**Public-source sample:** an unaffiliated demonstration from public evidence; the organization did not request, authorize, pay for, or participate in it.

**Buyer-requested outside-in review:** commissioned by a third party assessing the organization, on public and buyer-provided evidence.

## STANDARDS COMMITTEE ATTESTATION

This AI Trust File was produced and issued automatically under the Quiet AI Standard, on the basis of the evidence recorded herein. It carries no individual signature.

---

QUIET AI DIAGNOSTIC ENGINE · DIAZ INDUSTRIES INCORPORATED · STANDARD  
v3.1

# Suspension & disclaimer

---

## SUSPENSION OR REVOCATION

A certificate may be suspended or revoked if it is found to have been based on materially inaccurate information, or if circumstances change such that the recorded result no longer holds. Suspension removes any public verification record.

## DISCLAIMER AND LIMITATIONS

The certificate confirms a completed diagnostic at the purchased review depth and reports a current score; it does not mean the organization passed, is compliant, is safe, government-approved, ISO-certified, or free of AI risk. This file records an automated, evidence-based diagnostic review at a point in time, within the stated scope and review tier. It is not a legal opinion, a legal compliance certification, a government approval, an ISO certification, an insurance product, or a guarantee against regulatory action. It does not certify the safety or performance of any AI system.

### RELIANCE LIMITATION

This file may be relied on only as a point-in-time AI governance review within the stated scope and review tier.

### SAMPLE NOTICE

Illustrative sample. Built for a fictional organization to demonstrate the format. The organization did not participate and has not reviewed or approved this document; it is not an issued certificate. Figures reflect public evidence reviewed in June 2026 and are illustrative.

# Quiet AI™

PRIVATE STANDARDS BODY FOR AI GOVERNANCE

---

*Trust, made legible. Evidence over  
assertion, on every page.*