

Meridian Peak, Inc.

NYSE: MPKS

OBSERVED

OUTSIDE-IN

NOT CERTIFIED

RUN

mpks-illus-2026-09-29 · rev 3

ENGINE

AIQ™ V3.3.1 · Observed

EVALUATED

29 September 2026

EVIDENCE WINDOW

29 September 2025 – 29 September 2026

Supersedes mpks-sample-2026-09-29 · Analyst of record: n/a — illustrative build

FICTIONAL SUBJECT — ILLUSTRATIVE ONLY

OVERALL RATING

2.5

★★★★☆

Found 33 of 68 called-for disclosures, tier-weighted within each dimension and then weighted across the five.

CONFIDENCE GRADE

B 16 corroborated · 14 supported · 3 thin

33

FOUND

PROVENANCE
OF FOUND

33

NOT FOUND

High 24 — the company's own domain

2

LAPSED

Medium 6 — authored by the company, hosted elsewhere

17

SOURCES

Low 3 — third-party reporting only

DIMENSION	WEIGHT	CALLED FOR	FOUND	RATING	CONFIDENCE
● Strategic Alignment	20%	10	5	★★★★☆ 3.0	A
● Oversight & Accountability	30%	15	6	★★★★☆ 2.0	A
● Technical Robustness	25%	10	2	★★★★☆ 1.5	A
● Responsible AI & Compliance	15%	20	13	★★★★☆ 3.5	B
● Adaptability & Education	10%	13	7	★★★★☆ 2.0	B
Overall rating	100%	68	33	★★★★☆ 2.5	B

What the company has published

Strategy in public view. AI strategy, a two-year roadmap and measurable AI objectives are published on the company's own site and summarised in the 10-K.

Board oversight and staff rules. The proxy assigns AI oversight to the Audit & Risk Committee with semi-annual risk reports; a current acceptable-use policy bars client and personal data from public generative-AI tools.

Clients, data rights and training. A client AI-use statement is tied to the services terms; the assistant is labelled as AI; access-request, profiling-objection and transfer routes and a role-specific AI curriculum are published.

Where called-for disclosure is missing

Technical Robustness rates 1.5 stars. No independent validation, adversarial testing, testing methodology or documentation standard was found; human review of AI-drafted client work was last described June 2025 — lapsed.

Incident handling. No AI incident-response plan, incident tracking or post-incident review was found — holding Oversight & Accountability at 2.0 stars despite board oversight.

Management-layer controls. No ethics board, risk register or pre-deployment review was found; vendor-selection criteria last published November 2024 — lapsed; the sustainability report was unreachable and is recorded as not found.

Why the rating is defensible

Counter-verifier: one finding overturned, one clean record confirmed. A news profile described a pre-deployment review board no company source supports (recorded as not found); a check of enforcement and litigation records found no AI violations.

Correction and conservative reading. AI-accountability roles corrected from three to two after a departure; the lower of two conflicting steering-group readings was used.

Traceability. Every found disclosure cites the framework clause that calls for it and the public source where it was found; 24 of 33 rest on the company's own site, 3 on third-party sources only.

NOTICE

AIQ™ Observed Ratings and AIQ™ Scores are statements of opinion as of the date expressed, not statements of current or historical fact. Provided “as is” without warranties of any kind. An Observed Rating is built from public information without the subject company's participation and reflects only what that company has published. Where a called-for disclosure is absent, the rating records that it was not found — not that the underlying governance control is absent. Ratings do not constitute regulatory compliance, certification, or legal assurance. A named company may request review of the underlying public data.

For informational purposes only. Not investment, legal, tax, or compliance advice, and not a recommendation to buy, sell, or hold any security, or to do (or refrain from doing) business with any entity.

Fictional subject — illustrative only. Meridian Peak, Inc., its ticker, its sources, its peers and its findings are invented to show the format of an Observed report. Any resemblance to an actual company is coincidental.

S1 About AIQA

AIQA Global is an AI governance ratings firm. Its AIQ™ methodology rates how well an enterprise governs AI using 250 data points across five weighted dimensions: Strategic Alignment (20%), Oversight & Accountability (30%), Technical Robustness (25%), Responsible AI & Compliance (15%) and Adaptability & Education (10%). AIQA was built by the creators of the Ocean Tomo 300® Patent Index and applies the same discipline — a quality rating, an investable index, market adoption — to AI governance.

The AIQ™ Score is a 0–200 rating issued through a direct engagement, in which the company takes part and AIQA verifies the evidence. This report is an AIQ™ Observed Rating, a separate product. AIQA builds it from the outside, using only what the subject has published, without the company's participation and without access to its internal records. It measures how many of the disclosures the frameworks call for the company has made public, and it reports the result in stars rather than on the 0–200 scale. The two scales are not comparable, and one cannot be converted into the other.

S2 What an Observed Rating Is — and Is Not

The anchor frameworks — the NIST AI RMF, ISO/IEC 42001, the EU AI Act, the OECD AI Principles and GDPR — call on a company to publicly disclose specific things about how the organisation governs AI usage. An Observed Rating measures how many of those 68 framework-cited disclosures the company has actually published.

Six collectors gather evidence independently — SEC filings, the company's website, AI product screens, regulator records, patents and research, and ESG reports — and a counter-verifier tries to disprove every candidate finding. A finding counts only if it survives; where sources conflict, the more conservative reading wins.

Each disclosure carries a tier; results are weighted across the five dimensions and expressed on a scale of 0 to 5 in half-star steps, per dimension and overall. Every disclosure ends in one of three statuses: Found (published inside the 12-month evidence window), Not found (sources reviewed, disclosure not located) or Lapsed (published, but the latest evidence is more than 12 months old). Provenance — the company's own site, its filings, or third parties — is reported beside the rating.

What it is not. A missing disclosure is a missing disclosure — not proof that a control is missing. A company may govern AI well and publish little; it will rate lower than one that governs comparably and publishes what the frameworks call for. An Observed Rating is not an AIQ™ Score and not an AIQA Certification. A rated company may request a review of the public data behind its rating, and the disclosure-level evidence is available as an appendix under NDA. Engaging AIQA directly converts an Observed Rating into a verified AIQ™ Score.

S3 Analyst View

Meridian Peak publishes the top of its AI governance and very little of the machinery underneath it. Its AI strategy, roadmap and objectives are public. Its board has handed AI oversight to a committee that receives regular risk reports, employees work under a current acceptable-use policy, and clients are told how AI is used in their work.

What we could not find is the operating layer that would show an outsider whether those commitments hold up in practice: no AI incident-response plan, no risk register, no independent validation or adversarial testing, and no model documentation standard. Two disclosures that once existed — the criteria for choosing AI vendors and human review of AI-drafted client work — have not been refreshed in more than a year and now count as lapsed.

The result is a 2.5-star overall rating, strongest in Responsible AI & Compliance (3.5) and thinnest in Technical Robustness (1.5). We grade our confidence in this reading at B. Sixteen of the 33 credited disclosures are corroborated: they appear in more than one place or more than one reporting year, and the sources agree. The overall rating does not lean on third-party evidence — removing the three third-party-only credits leaves it unchanged, though Responsible AI & Compliance would drop half a star.

§4 Scope of This Run

EVALUATED	29 September 2026 · evidence window 29 September 2025 – 29 September 2026
RETRIEVED	24 September 2026 · Recurrence fields added 4 October 2026 (rev 1); statuses, credit, provenance and ratings unchanged.
CALLED FOR	68 — Strategic Alignment 10 · Oversight & Accountability 15 · Technical Robustness 10 · Responsible AI & Compliance 20 · Adaptability & Education 13
STATUS	Found 33 · Not found 33 · Lapsed 2
PROVENANCE	High 24 · Medium 6 · Low 3 (found disclosures)
CONFIDENCE GRADE	B · disclosures graded A 16 · B 14 · C 3 · by dimension A / A / A / B / B
SOURCES	17 distinct public sources inside the window; 6 earlier editions consulted for recurrence only (no credit)
COLLECTORS	SEC filings (proxy, 10-K, 8-K) and the company website (11 pages) returned the credited evidence; third-party news supplied two credits, and the counter-verifier’s check of regulator and court records supplied one. AI product screens, patents & research and ESG reports returned none — the sustainability report could not be reached on two attempts.

Where the credit came from

Each found disclosure traces to a document Meridian Peak or a third party published. A disclosure seen in two places counts under both; earlier editions count toward recurrence only and earn no credit.

CODE	SOURCE	LOCATION (FICTIONAL)	PROV.	DATED	FOUND DISCLOSURES CARRIED
S13	Form 10-K, FY2025	edgar.example/MPKS/10-K-FY2025	Medium	20 Feb 26	6 — #3, #6, #14, #36, #37, #51
S01	Responsible-AI page	meridianpeak.example/responsible-ai	High	18 Jun 26	5 — #1, #20, #40, #41, #65
S04	Privacy notice	meridianpeak.example/privacy	High	15 Jan 26	5 — #27, #45, #47, #48, #49
S02	AI acceptable-use policy	meridianpeak.example/trust/ai-acceptable-use	High	3 Mar 26	4 — #14, #15, #24, #29
S03	Investor AI strategy page	meridianpeak.example/investors/ai-strategy	High	14 May 26	3 — #4, #5, #6
S06	Client AI-use statement	meridianpeak.example/clients/ai-use-statement	High	1 Jul 26	3 — #29, #41, #42
S12	Proxy statement 2026 (DEF 14A)	edgar.example/MPKS/DEF14A-2026	Medium	10 Apr 26	3 — #1, #11, #13
S14	Form 8-K, May 2026 investor day	edgar.example/MPKS/8-K-2026-05-investor-day	Medium	14 May 26	3 — #4, #60, #66
S07	Assistant FAQ	meridianpeak.example/support/assistant-faq	High	2 May 26	2 — #42, #61
S08	Careers: learning	meridianpeak.example/careers/learning	High	22 Apr 26	2 — #37, #53
S11	Feedback page	meridianpeak.example/feedback	High	2 May 26	2 — #61, #63
S05	Privacy request portal	meridianpeak.example/privacy/requests	High	15 Jan 26	1 — #48
S09	Innovation-lab page	meridianpeak.example/innovation-lab	High	20 May 26	1 — #58
S10	Company newsroom	meridianpeak.example/newsroom	High	5 Aug 26	1 — #60
S15	News report: AI partnerships	news.example/meridian-peak-ai	Low	11 Feb 26	1 — #59
S16	News report: state AI hearing	news.example/state-ai-hearing-2026	Low	19 Mar 26	1 — #46
S17	Clean-record check: enforcement and litigation records	records.example/cv-clean-record-check-mpks	Low	24 Sep 26	1 — #44
P1	Proxy statement 2025 (DEF 14A)	edgar.example/MPKS/DEF14A-2025	Medium	11 Apr 25	recurrence — #1, #11
P2	Form 10-K, FY2024	edgar.example/MPKS/10-K-FY2024	Medium	21 Feb 25	recurrence — #3, #6
P3	AI acceptable-use policy, 2025 edition	meridianpeak.example/trust/ai-acceptable-use?ed=2025	High	4 Mar 25	recurrence — #15
P4	Privacy notice, 2025 edition	meridianpeak.example/privacy?ed=2025	High	16 Jan 25	recurrence — #27
P5	Responsible-AI page, 2025 edition	meridianpeak.example/responsible-ai?ed=2025	High	20 Jun 25	recurrence — #40
P6	Careers: learning, 2025 edition	meridianpeak.example/careers/learning?ed=2025	High	15 Apr 25	recurrence — #53

All sources retrieved 24 September 2026. P-rows are earlier editions outside the evidence window. The 10-K sits behind six found disclosures and the privacy notice and responsible-AI page behind five each. Every credit but three rests on something Meridian Peak wrote itself; 24 of the 33 rest on its own website, which makes the record easy to check — and easy to change, which is why the lapsed and corrected items matter.

\$5 Strategic Alignment

★★★★★ 3.0 Weight 20% · Found 5 of 10 called-for · Not found 4 · Lapsed 1 · Confidence Grade A

What it measures. Whether AI is tied to the business plan: who owns it at the top, whether the strategy is written down with goals and a roadmap, what the company invests, and how it picks AI suppliers.

What the frameworks call for. The OECD AI Principles and NIST AI RMF (GOVERN) call for a public strategy and executive ownership; ISO/IEC 42001 (clauses 4-6) adds a context analysis, measurable objectives and a roadmap. Investment, resourcing and vendor-selection disclosures make up the rest.

What we found for Meridian Peak, Inc.

The strategy is fully in public view. The investor AI page sets out the strategy, a two-year roadmap and three measurable objectives, and the 10-K repeats the strategy and says where AI is used.

The proxy describes an executive AI steering group; we credited it in part and took the more conservative of two readings, because the proxy and the website describe its meeting cadence differently. Nothing was found on AI spending, budget share, resource adequacy or context analysis, and the vendor-selection criteria last appeared in November 2024.

A strong dimension and the cheapest to improve: republishing the vendor criteria alone would lift it from 3.0 to 4.0 stars (illustrative).

PUBLISHED	CALLED FOR, NOT FOUND
— AI strategy, on the investor page and in the 10-K — Two-year AI roadmap — Three measurable AI objectives — AI use described in SEC filings — Executive AI steering group (partial)	— Business-context analysis for AI — AI spending relative to revenue — AI share of the IT budget — Assessment of AI resource adequacy

Lapsed: AI vendor-selection criteria — last evidence 2 Nov 2024 (responsible-AI page).

Where each credit came from

FOUND DISCLOSURE	WHERE FOUND (SOURCE, DATE) · EARLIER EDITIONS	PLACES · PERIODS	PROV.	CONF.
#1 Executive AI steering group (partial)	S12 Proxy 2026 (10 Apr 26); S01* Responsible-AI page, offline capture (18 Jun 26) · earlier: P1 Proxy 2025	2 · 2	High*	B · conservative reading
#3 AI described in SEC filings	S13 10-K FY2025 (20 Feb 26) · earlier: P2 10-K FY2024	1 · 2	Medium	A
#4 Three measurable AI objectives	S03 Investor AI page (14 May 26); S14 8-K, investor day (14 May 26)	2 · 1	High	A
#5 Two-year AI roadmap	S03 Investor AI page (14 May 26)	1 · 1	High	B · single source
#6 Public AI strategy	S03 Investor AI page (14 May 26); S13 10-K FY2025 (20 Feb 26) · earlier: P2 10-K FY2024	2 · 2	High	A

S6 Oversight & Accountability

★ ★ ★ ★ ★ 2.0 Weight 30% · Found 6 of 15 called-for · Not found 9 · Lapsed 0 · Confidence Grade A

What it measures. Who answers for AI and how problems get caught: board oversight, written policy, named accountability, audit and monitoring, and risk management. It carries the most weight of the five.

What the frameworks call for. NIST AI RMF (GOVERN) and ISO/IEC 42001 (clauses 5, 6.1 and 10) call for board oversight, a current policy, named roles, an ethics or review body, an incident-response plan, third-party AI oversight, pre-deployment review and a risk register; the EU AI Act (Arts. 9 and 17) asks the same of high-risk providers.

What we found for Meridian Peak, Inc.

Oversight at the top is visible. The proxy assigns AI oversight to the Audit & Risk Committee, which receives semi-annual AI risk reports from the Chief Operating Officer. The acceptable-use policy is current (March 2026), screens proposed uses against a prohibited list and is referenced in the 10-K. The responsible-AI page names two accountable roles — corrected from three after one executive left in July 2026.

Below that, the record goes quiet: no incident-response plan, ethics body, risk register, third-party oversight, board training or quality system. A news profile described a pre-deployment review board, but no company source supports it, so the counter-verifier overturned it.

A board that is told about AI risk, with no public evidence of the process that produces what it is told. The missing incident-response plan is the single most material gap in this report.

PUBLISHED — Board committee assigned AI oversight — Semi-annual AI risk reports to the board — Current AI acceptable-use policy (Mar 2026) — Screening against prohibited AI uses — Two named AI-accountable roles (partial)	CALLED FOR, NOT FOUND — AI incident-response plan — AI risk register — AI ethics or review body — Pre-deployment review (counter-verifier overturned) — Oversight of third-party AI — Board AI training
---	--

Where each credit came from

FOUND DISCLOSURE	WHERE FOUND (SOURCE, DATE) · EARLIER EDITIONS	PLACES · PERIODS	PROV.	CONF.
#11 Board committee oversees AI	S12 Proxy 2026 (10 Apr 26) · earlier: P1 Proxy 2025	1 · 2	Medium	A
#13 Semi-annual AI risk reports to the board	S12 Proxy 2026 (10 Apr 26)	1 · 1	Medium	B · single source
#14 AI acceptable-use policy	S02 AI acceptable-use policy (3 Mar 26); S13 10-K FY2025 (20 Feb 26)	2 · 1	High	A
#15 AI policy updated within the year	S02 AI acceptable-use policy (3 Mar 26) · earlier: P3 AUP, 2025 ed.	1 · 2	High	A
#20 Two named AI-accountable roles (partial)	S01 Responsible-AI page (18 Jun 26)	1 · 1	High	B · corrected
#24 Screening against prohibited AI uses	S02 AI acceptable-use policy (3 Mar 26)	1 · 1	High	B · single source

S7 Technical Robustness

★ ★ ★ ★ ★ 1.5 Weight 25% · Found 2 of 10 called-for · Not found 7 · Lapsed 1 · Confidence Grade A

What it measures. Whether AI systems are tested, secured, documented and checked before their output is used — the evidence that the technology works as intended.

What the frameworks call for. NIST AI RMF (MEASURE, MANAGE), ISO/IEC 42001 Annex A and the EU AI Act (Arts. 10, 11 and 15) call for independent validation, adversarial testing, a testing method, a documentation standard such as model cards, retention and erasure controls, recovery capability and human review of output.

What we found for Meridian Peak, Inc.

Only the data rules are published: the privacy notice’s retention schedule covers data processed by AI tools, and the acceptable-use policy — echoed in the client AI-use statement — bars client confidential and personal data from public generative-AI tools.

Human review of AI-drafted client deliverables was described in June 2025 and not since. No validation, adversarial testing, testing method, documentation standard, erasure capability, recovery figure or AI-specific security-incident reporting was found.

For a services firm whose AI drafts client work, the lapsed output review matters most. Republishing it would lift this dimension to 2.0 stars; adding independent validation and a documentation standard would take it from 1.5 to 2.5 (illustrative).

PUBLISHED	CALLED FOR, NOT FOUND
— Retention policy covering AI data — Data types barred from AI tools	— Independent validation of AI systems — Adversarial or red-team testing — Documented testing and evaluation method — Model documentation standard — Training-data erasure capability — AI security-incident reporting

Lapsed: Human review of AI-drafted client deliverables — last evidence 30 Jun 2025 (client AI-use statement).

Where each credit came from

FOUND DISCLOSURE	WHERE FOUND (SOURCE, DATE) · EARLIER EDITIONS	PLACES · PERIODS	PROV.	CONF.
#27 Retention policy covers AI data	S04 Privacy notice (15 Jan 26) · earlier: P4 Privacy notice, 2025 ed.	1 · 2	High	A
#29 Data types barred from AI tools	S02 AI acceptable-use policy (3 Mar 26); S06 Client AI-use statement (1 Jul 26)	2 · 1	High	A

S10.1 Verification record

CHECK	N	DISCLOSURE AND DISPOSITION
Counter-verifier overturned	1	#23 Pre-deployment review — a news profile described a review board; no company source supports it. Recorded as not found.
Conservative reading	1	#1 Executive AI steering group — the proxy and the company website describe meeting cadence differently; the lower reading was used (partial credit).
Corrected	1	#20 Named AI-accountable roles — corrected from three to two after a named executive left in July 2026 and was removed from the site.
Lapsed	2	#10 Vendor-selection criteria (last evidence 2 Nov 2024); #35 Human review of AI-drafted client work (30 Jun 2025). No credit; reported as lapsed, not absent.
Source unreachable	1	#50 Environmental footprint — the sustainability report returned an error on collection and re-collection. Recorded as not found.
Clean-record check	1	#44 Regulatory violations — the company reports nothing on AI violations. The counter-verifier checked enforcement and litigation records (EU AI Office and national authority notices, AI litigation trackers, court dockets, FTC and state AG actions) and found none. Credited from third-party records (provenance Low).
Subsidiary attribution	0	None. Every credit is attributable to Meridian Peak, Inc. itself.

Two citations stand behind every one of the 33 found disclosures: the framework clause that calls for it and the public source where it was found, with source and retrieval dates (evidence appendix, available on request under NDA; source key above).

S8 Responsible AI & Compliance

★★★★★ 3.5 Weight 15% · Found 13 of 20 called-for · Not found 7 · Lapsed 0 · Confidence Grade B

What it measures. How the company treats the people its AI touches — employees, clients and the public — through transparency, rights and recourse, training, compliance, and environmental and supply-chain responsibility.

What the frameworks call for. GDPR (Arts. 15, 21, 22) and the EU AI Act (Arts. 4, 50, 86) call for access and objection rights, notice when people deal with AI, AI literacy and a route to contest AI decisions; the OECD AI Principles and ISO/IEC 42001 Annex A add transparency reporting, workforce planning and supply-chain diligence.

What we found for Meridian Peak, Inc.

This is Meridian Peak’s strongest dimension: 13 of 20 found. Clients get an AI-use statement that the services terms reference; the customer-service assistant identifies itself as AI and hands off to a person; the privacy site covers access requests, profiling objections and the legal basis for cross-border AI data; and the careers site describes AI training, including a curriculum built for schedulers and case handlers. The 10-K reports AI literacy coverage and ethics-training completion.

One partial credit deserves plain words: the recourse route is request-based — people can appeal and ask for human review, but no resolution tracking or response times are published. Regulator engagement rests on one news report. Meridian Peak reports nothing on AI regulatory violations, so the counter-verifier checked enforcement and litigation records and confirmed a clean record; that disclosure is credited from third-party records. The company has no published plan for reporting a violation if one occurs, and we would encourage one. No transparency report, workforce or communication plan, or supply-chain diligence was found, and the unreachable sustainability report leaves the environmental disclosure not found.

Strong on individual rights and client notice; thin on reporting outward.

PUBLISHED	CALLED FOR, NOT FOUND
— Client AI-use statement, referenced in services terms — AI labelling with human handoff — Access, objection and cross-border transfer rights — AI training, including a role-specific curriculum — AI literacy and ethics-training figures in the 10-K — Clean regulatory record (counter-verifier confirmed)	— AI transparency report — Workforce transition plan — AI communication plan — Supply-chain labour diligence — Baseline AI literacy assessment — Environmental footprint (source unreachable)

Where each credit came from

FOUND DISCLOSURE	WHERE FOUND (SOURCE, DATE) · EARLIER EDITIONS	PLACES · PERIODS	PROV.	CONF.
#36 AI literacy coverage reported	S13 10-K FY2025 (20 Feb 26)	1 · 1	Medium	B · single source
#37 AI training programmes for staff	S08 Careers: learning (22 Apr 26); S13 10-K FY2025 (20 Feb 26)	2 · 1	High	A
#40 Responsible-AI disclosures published	S01 Responsible-AI page (18 Jun 26) · earlier: P5 Responsible-AI page, 2025 ed.	1 · 2	High	A
#41 Clients told how AI is used	S06 Client AI-use statement (1 Jul 26); S01 Responsible-AI page (18 Jun 26)	2 · 1	High	A
#42 Assistant identifies itself as AI	S07 Assistant FAQ (2 May 26); S06 Client AI-use statement (1 Jul 26)	2 · 1	High	A
#44 Clean regulatory record (counter-verifier confirmed)	S17 Clean-record check: enforcement & litigation records (24 Sep 26)	1 · 1	Low	C · third-party only
#45 Legal basis for cross-border AI data	S04 Privacy notice (15 Jan 26)	1 · 1	High	B · single source
#46 Engagement with AI regulators	S16 News: state AI hearing (19 Mar 26)	1 · 1	Low	C · third-party only
#47 Recourse route (partial)	S04 Privacy notice (15 Jan 26)	1 · 1	High	B · single source
#48 Access requests cover AI data	S05 Privacy request portal (15 Jan 26); S04 Privacy notice (15 Jan 26)	2 · 1	High	A
#49 Objection to AI profiling	S04 Privacy notice (15 Jan 26)	1 · 1	High	B · single source
#51 AI ethics training completion (partial)	S13 10-K FY2025 (20 Feb 26)	1 · 1	Medium	B · single source
#53 Role-specific AI curriculum	S08 Careers: learning (22 Apr 26) · earlier: P6 Careers, 2025 ed.	1 · 2	High	A

\$9 Adaptability & Education

★ ★ ★ ★ ★ 2.0 Weight 10% · Found 7 of 13 called-for · Not found 6 · Lapsed 0 · Confidence Grade B

What it measures. Whether the company learns and adjusts: how it tracks and reviews AI incidents, listens to outsiders, manages change, experiments and measures improvement.

What the frameworks call for. NIST AI RMF (MANAGE 4) and ISO/IEC 42001 (clauses 9 and 10) call for incident tracking and post-incident review, feedback channels with a set way to resolve what comes in, monitoring and change management; the EU AI Act (Art. 57) adds regulatory sandboxes.

What we found for Meridian Peak, Inc.

Meridian Peak shows that it experiments and listens. An innovation-lab page lists AI pilots; the newsroom and investor-day 8-K describe AI features released this year; a feedback page offers three channels, one through the assistant; clients and employees have a stated route to comment on AI changes; a change-management process covers AI rollouts; and the 8-K cites an external AI maturity assessment. AI partnerships are credited from news coverage only.

Not found: incident tracking, post-incident review, a process for resolving feedback, real-time monitoring, measurement of whether changes work, or sandbox participation.

Inputs without the loop that closes them — and with no incident-response plan in Oversight either, the incident gap runs across two dimensions.

PUBLISHED	CALLED FOR, NOT FOUND
— Experimental AI projects (innovation lab) — New AI capabilities released in the past year — Three external feedback channels (partial) — Stakeholder input on AI changes — Change management for AI rollouts — External AI maturity assessment	— AI incident tracking — Post-incident reviews — Process to resolve external feedback — Real-time AI performance monitoring — Measurement of improvement effectiveness — Regulatory sandbox participation

Where each credit came from

FOUND DISCLOSURE	WHERE FOUND (SOURCE, DATE) · EARLIER EDITIONS	PLACES · PERIODS	PROV.	CONF.
#58 Experimental AI projects	S09 Innovation-lab page (20 May 26)	1 · 1	High	B · single source
#59 External AI partnerships	S15 News: AI partnerships (11 Feb 26)	1 · 1	Low	C · third-party only
#60 New AI capabilities in the past year	S10 Newsroom (5 Aug 26); S14 8-K, investor day (14 May 26)	2 · 1	High	A
#61 External feedback channels (partial)	S11 Feedback page (2 May 26); S07 Assistant FAQ (2 May 26)	2 · 1	High	A
#63 Stakeholder input on AI changes	S11 Feedback page (2 May 26)	1 · 1	High	B · single source
#65 Change management for AI rollouts	S01 Responsible-AI page (18 Jun 26)	1 · 1	High	B · single source
#66 External AI maturity assessment	S14 8-K, investor day (14 May 26)	1 · 1	Medium	B · single source

\$10.2 Confidence Grade B

Confidence Grade. How sure we are that this reading is reliable. The stars measure how much the company has published. The Confidence Grade measures the evidence behind each credit: whether it appears in more than one place, whether it has appeared in more than one reporting period, whether the sources agree, and whether at least one source is the company’s own. It is reported beside the rating and never changes a star.

GRADE	RULE (PER FOUND DISCLOSURE)	MERIDIAN PEAK, INC.
A • Corroborated	Company-authored source; appears in two or more in-window sources or two or more reporting periods; every source agrees.	16
B • Supported	Company-authored, but one source in one period (not tested against another), or sources disagreed and the conflict was resolved by correction or the conservative reading.	14
C • Thin	Third-party sources only. Still credited — provenance and confidence never remove credit.	3

Run level. A when at least two-thirds of found disclosures grade A and all six ratings hold with every grade-C credit removed; B when at least one-third grade A and the overall rating holds without the grade-C credits; C otherwise. Meridian Peak has 16 of 33 at grade A, just under half, where two-thirds would need 22. Without its three grade-C credits the overall still holds, but Responsible AI & Compliance would fall from 3.5 to 3.0. Result: B.

Dimension level. Each found disclosure scores A = 2, B = 1, C = 0. A dimension’s grade is the average of its disclosures: A at 1.5 or above, B from 0.5 up to 1.5, C below 0.5. The overall grade follows the run-level rule above. For Meridian Peak, Inc. the overall average (1.39) points the same way.

DIMENSION	RATING	FOUND	A	B	C	AVERAGE	GRADE	HOLDS WITHOUT C
Strategic Alignment	3.0	5	3	2	0	1.60	A	Yes
Oversight & Accountability	2.0	6	3	3	0	1.50	A	Yes
Technical Robustness	1.5	2	2	0	0	2.00	A	Yes
Responsible AI & Compliance	3.5	13	6	5	2	1.31	B	No
Adaptability & Education	2.0	7	2	4	1	1.14	B	Yes
Overall	2.5	33	16	14	3	1.39	B	Yes

Margin. The grade measures the evidence, not how close the rating sits to a boundary — so we report margin separately. Meridian Peak’s overall sits just above the line between 2.0 and 2.5 stars. Losing any one of ten credits on refresh would move it to 2.0 (illustrative): seven of them grade A, and three grade B — the executive AI steering group (#1, conservative reading), the AI roadmap (#5) and the board’s AI risk reports (#13), each resting on a single document. Those three are the first to re-check.

What would raise the grade. Six more corroborated disclosures would meet the share test for A, but A also needs every rating to hold without third-party-only credits, and Responsible AI & Compliance’s half-star now rests on one. Part of the path needs nothing new from Meridian Peak: if the 2027 proxy repeats the board’s AI risk reporting and the next 10-K repeats the roadmap and the literacy and ethics-training figures, those four credits move from B to A on the next annual cycle. Two more corroborations would complete the move.

Low provenance still counts. Under the published methodology, provenance and confidence sit beside the rating; neither removes credit. The three grade-C credits — the clean regulatory record (#44), regulator engagement (#46) and AI partnerships (#59) — are in the stars. The overall does not depend on them; Responsible AI & Compliance’s move to 3.5 does.

§11 Summary Assessment

Meridian Peak reads strongest in Responsible AI & Compliance (3.5) and Strategic Alignment (3.0) — how it treats the people AI touches and what it intends to do with AI. It reads thinnest in Technical Robustness (1.5) and, at 2.0, in Oversight & Accountability and Adaptability & Education: the evidence that its controls work and that it learns when they fail.

A simple average of the five dimension ratings would be 2.4. The overall is 2.5 because it is computed from the weighted, unrounded dimension results rather than from the stars, so it does not inherit each dimension’s rounding.

§11.1 What earned no credit

REASON	N	DISCLOSURES
Not found — reviewed, not located	31	The most material among them: board AI training; AI ethics or review body; AI incident-response plan; oversight of third-party AI; AI risk register; independent validation; AI security-incident reporting; model documentation standard; AI literacy baseline; supply-chain labour diligence; AI incident tracking; post-incident reviews. The other 19 are listed in the evidence appendix, available on request under NDA.
Not found — counter-verifier overturned	1	#23 Pre-deployment review
Not found — source unreachable	1	#50 Environmental footprint governance
Lapsed — evidence older than 12 months	2	#10 Vendor-selection criteria (Nov 2024); #35 Human review of AI-drafted client work (Jun 2025)
Credited in part	5	#1 Steering group (conservative reading); #20 Accountable roles (corrected); #47 Recourse route; #51 Ethics-training completion; #61 Feedback channels
Credited on third-party evidence only	3	#44 Clean regulatory record (counter-verifier check); #46 Regulator engagement; #59 AI partnerships — credited, grade C. The overall does not depend on them; Responsible AI & Compliance’s 3.5 depends on #44

35 of the 68 earned no credit (Not found 33 + Lapsed 2); 8 more earned credit with a qualification.

§11.2 Dependencies and limits

- No credit depends on a subsidiary. Provenance is concentrated on the company’s own website — easy to verify, and easy to edit or withdraw without notice.
- The rating supports a view of what Meridian Peak has chosen to publish about AI governance as of 29 September 2026, comparable with any other rating on the same 68 disclosures.
- It does not support a conclusion that the controls we did not find are absent, a view of how its AI systems perform, or any investment, legal or compliance judgment.

§11.3 Forward look (illustrative)

- Refreshing the two lapsed disclosures would lift Strategic Alignment from 3.0 to 4.0 stars and Technical Robustness from 1.5 to 2.0.
- Adding an AI incident-response plan would lift Oversight & Accountability from 2.0 to 2.5; together with the two refreshes, the overall would move from 2.5 to 3.0.
- Publishing an incident-response plan, incident tracking and post-incident reviews would lift Oversight & Accountability to 2.5 and Adaptability & Education from 2.0 to 3.5.
- Publishing independent validation and a model documentation standard would lift Technical Robustness from 1.5 to 2.5.
- On the downside, losing any of the ten margin credits would take the overall to 2.0.

\$12 Comparative Context

No peer runs exist yet on the 68 disclosures, so this section uses Mode A: recommended peers, without ratings. Every company below is fictional.

COMPANY	AXIS	AI-GOVERNANCE DISCLOSURE POSTURE	WHY A PEER
Peer A (fictional) — outsourced scheduling and case management	Business services · enterprise AI use	Responsible-AI page and acceptable-use summary; board AI oversight in its proxy	Same client-delivery model and AI use in scheduling
Peer B (fictional) — claims and document processing	Business services · enterprise AI use	Discusses AI in its 10-K; little on its own website	Similar document workload; tests a filings-only disclosure pattern
Peer C (fictional) — workforce solutions	Professional services · enterprise AI use	Publishes an AI transparency report; no board-level AI oversight disclosed	Contrast: strong outward reporting, thin oversight
Peer D (fictional) — customer-contact outsourcing	Business services · enterprise AI use	Assistant disclosures and model cards for client-facing tools	Same assistant exposure; shows documentation Meridian Peak lacks

Caveat. Companies differ widely in how much they publish about AI governance, and a star rating measures publication, not practice. A company may govern AI well and disclose little; it will rate lower than one that governs comparably and publishes what the frameworks call for. Compare within an industry axis before drawing conclusions across axes. Mode B (comparative ratings) requires peer runs on the same 68 disclosures.

\$13 Definitions

Called-for disclosure	One of the 68 items the anchor frameworks call on a company to publish about AI governance. The other 182 of the 250 base data points are not publicly disclosable by design.
Found / Not found / Lapsed	Found: located in a public source dated inside the 12-month window, attributable to the company, and not overturned — or, for regulatory violations, a clean record the counter-verifier confirmed when the company reports nothing. Not found: sources reviewed, disclosure not located. Lapsed: located, but the latest evidence is older than 12 months. Only Found earns credit.
Provenance	High: the company's own website or domain. Medium: company-authored content hosted elsewhere, such as an SEC filing. Low: third-party commentary. The strongest available source applies. An asterisk (High*, Medium*, Low*) means AIQA relied on an offline HTML capture rather than the live page; the copy may display inaccurately or be corrupted. The asterisk never changes the level, which follows where the content came from. Reported beside the rating; never applied to it.
Confidence Grade	A Corroborated · B Supported · C Thin for each found disclosure. A dimension's grade is the average of its disclosures (A = 2, B = 1, C = 0; A at 1.5 or above, B at 0.5 or above); the overall grade follows the run-level rule. Built from recurrence across places and reporting periods, agreement between sources, and company authorship. Reported beside the rating; never applied to it.
Earlier edition	A prior-year filing or earlier version of a company page (P1-P6). Counts toward recurrence only; it sits outside the window and earns no credit.
Star scale	Ratings run from 0 to 5 in half-star steps, for each dimension and overall. The overall comes from the weighted dimension results, not an average of stars.

\$14 Disclosure-Level Evidence

Disclosure-Level Evidence lists all 68 called-for disclosures with the proof behind each one: question, framework clause, status, provenance, recurrence, Confidence Grade and source. It is not part of this report and is available for audit on request, under a non-disclosure agreement.

RUN INTEGRITY

Run mpks-illus-2026-09-29 · rev 3 · Evaluated 29 September 2026 · AIQ™ V3.3.1 · Observed

Findings workbook CONFIDENTIAL_AIQ_Observed_MPKS_findings_illus_rev3.xlsx

SHA-256 970ea33b710353b94e42483a0738b7c0b8cc248e865affbd05b5621acbf5470

NOTICE

AIQ™ Observed Ratings and AIQ™ Scores are statements of opinion as of the date expressed, not statements of current or historical fact. Provided “as is” without warranties of any kind. An Observed Rating is built from public information without the subject company's participation and reflects only what that company has published. Where a called-for disclosure is absent, the rating records that it was not found — not that the underlying governance control is absent. Ratings do not constitute regulatory compliance, certification, or legal assurance. A named company may request review of the underlying public data.

For informational purposes only. Not investment, legal, tax, or compliance advice, and not a recommendation to buy, sell, or hold any security, or to do (or refrain from doing) business with any entity.

Fictional subject — illustrative only. Meridian Peak, Inc., its ticker, its sources, its peers and its findings are invented to show the format of an Observed report. Any resemblance to an actual company is coincidental.