AI Board Readiness Checklist

12 Questions Before Backing an AI Company or Approving an AI Mandate

Prepared for Non-Executive Directors, Board Chairs, Venture Partners, and Family Office Principals. Grounded in operating multi-agent production workflows, building verifiable AI compute, and analysing primary SEC AI filings at Venture Signal.

PART I · CAPITAL, UNIT ECONOMICS & DEFENSIBILITY

INVESTMENT & STRATEGY COMMITTEES

01 What happens to gross margin if top-decile usage triples under flat pricing?

☐ AI carries variable inference cost per request. Benchmark data places average AI product gross margins around 45% to 53% versus 65% to 70% for traditional SaaS, and flat monthly subscriptions invert rapidly when heavy customers scale token consumption.

RED FLAG Blended gross margin presented without a decile breakdown of token cost per active seat.

02 Does the architecture route by task complexity, or pay frontier rates for routine work?

☐ In production pipelines, switching between frontier and lightweight model tiers alters long-document inference cost up to eightfold. Relying on endless prompt tuning inside one flagship API masks structural cost waste.

RED FLAG Zero multi-model routing policy and no benchmark showing where cheaper or local models replace frontier calls.

03 If the foundation model vendor doubles prices tomorrow, where does the moat live?

☐ Renting intelligence from a frontier provider is viable only when durable defensibility sits elsewhere: proprietary workflow state, regulated distribution, exclusive ground-truth data, or physical integration.

RED FLAG Switching away from the primary LLM leaves a thin prompt wrapper with minimal customer switching costs.

04 How much apparent demand depends on vendor credits or non-cancellable compute bets?

☐ Primary-filing analysis shows infrastructure suppliers extending hundreds of billions in purchase commitments, cloud buy-back pledges, and credit guarantees. Target companies may look profitable only while subsidised credits hold.

RED FLAG Gross margins calculated using promotional cloud credits rather than unsubsidised commercial unit rates.

PART II · OPERATIONAL GOVERNANCE & PRODUCTION RISK

AUDIT & RISK COMMITTEES

05 Who is accountable at 3 a.m. when an automated system ships a confident error?

☐ Boards frequently debate model selection while leaving overnight operational ownership undefined. Every production system requires a named owner, automated anomaly alerting, and explicit authority to pause execution immediately.

RED FLAG “The AI team monitors it” with no written escalation SLA or emergency kill-switch owner.

06 Does the system verify live state changes, or confuse an HTTP 200 status with reality?

☐ In multi-agent production logs, agents routinely treat API acknowledgements or queued job IDs as completed business outcomes. Execution must be decoupled from a separate check inspecting the resulting artefact.

RED FLAG Workflows that mark tasks complete solely from return codes without post-action state verification.

07 **What stops an external PDF, email, or web page from hijacking the agent's instructions?**

☐ Routine document screening exposes language models to untrusted text containing hidden instructions. Treating retrieved data as executable instruction is an architectural vulnerability, not a policy issue solved by staff guidance.

RED FLAG Unsandboxed tool calls triggered directly from external document content without instruction-data separation.

08 **Are bulk edits, external communications, and deploys behind human confirmation gates?**

☐ Unbounded autonomous loops turn minor logic errors into portfolio-wide damage within seconds. Scoped permissions, bulk-size circuit breakers, and explicit human confirmation gates prevent silent mass overwrites.

RED FLAG Production write credentials or outbound messaging permissions granted to autonomous workers without batch limits.

PART III · VERIFICATION, PROVENANCE & ACCOUNTABILITY

COMPLIANCE & NOMINATION

09 **Can an independent verifier check the output of a stronger model before customers see it?**

☐ Deterministic citation checks and secondary verification passes catch unverified or invented references before release. Boards should know which error classes a lightweight verifier catches reliably and where verification costs exceed direct execution.

RED FLAG Relying on the same generative pass to both author and self-certify factual or legal outputs.

10 **Is compute or output provenance attested cryptographically, or accepted on vendor trust?**

☐ When workloads run across third-party clouds or heterogeneous nodes, self-reported dashboards are weak evidence. Cryptographic signing of output artefacts shifts verification from trusting hardware claims to checking proof.

RED FLAG Regulated or external compute pipelines unable to prove which model version produced a given decision.

11 **Does management cap automated throughput where verification is the intentional bottleneck?**

☐ Disciplined operators keep batch sizes bounded so review capacity matches output volume, and distinguish prototype telemetry from audited commercial scale rather than inflating vanity usage metrics.

RED FLAG Uncapped AI output targets rewarded ahead of defect rates, rollback counts, or citation accuracy.

12 **If a regulator, buyer, or LP queries an AI claim, can the board produce primary proof in 48 hours?**

☐ Institutional trust rests on traceability: retained input snapshots, model routing logs, failed-batch registers, and clear boundary statements documenting what a system cannot yet verify automatically.

RED FLAG Ad-hoc prompt histories with no retained audit trail tying published outputs to primary sources.

BOARD SCORING & PRE-DEAL TRIAGE GUIDE

10 – 12 Clear Answers	6 – 9 Partial Answers	Under 6 Clear Answers
Prepared for capital allocation or production mandate under standard quarterly committee oversight.	Proceed only with milestone covenants: margin-by-decile reporting, confirmation gates, and failure logging.	Narrative is outpacing engineering governance. Pause capital or mandate pending an independent operator read.

George Bregman

AI Strategist, Venture Builder & Editor-in-Chief, Venture Signal (London & Cambridge, UK). CEO of AI Disraeli; Architect of Aurora Protocol; independent pre-deal second opinions for private capital. Holds one non-executive seat at a time and two to three advisory mandates per quarter.

georgebregman.com/board
george@georgebregman.com
venturesignal.org
linkedin.com/in/george-bregman

© 2026 George Bregman · General governance framework, not legal or investment advice