Authority, Cost & Control in AI Systems

Building on Vol. 2026.04, which defines Authority Architecture, this Control Note translates those principles into **Board-level oversight and enforceable controls**.

Signal from Market Reality

- AI systems do not retain operational memory; context and reasoning are reconstructed each time.
- What is called “AI memory” (RAG, context windows, embeddings) enables reconstruction not retention.
- As a result, **cognition is externalised and quantified** into tokens: reasoning, validation, and correction are paid for repeatedly.
- Scaling AI increases computation through repeated validation, not just usage.
- Reported growth (e.g., token volume) may reflect reconstruction overhead, not proportional value.

Structural Implication

- Human interactions reduce cost through memory and shared context or cognitive load (compression).
- AI systems maintain or increase cost through repeated reconstruction (expansion).
- **The burden of remembering** shifts from the AI to the user and is incurred as computation (tokens).

Control for Board Reframing

- Cost is driven by **repeated decision reconstruction**, not only model usage.
- Undefined authority leads to **unbounded computation and cost escalation.**
- Governance layers can **multiply cost** if not explicitly designed.
- Value cannot be inferred from activity; it must be measured at **decision level.**

Enforcement

- **Board Risk Register** defines structural exposures arising from reconstruction and authority ambiguity.
- **Board Audit Checklist** defines the non-negotiable conditions required to contain cost, control execution, and ensure accountability.

Autonomy scales. Authority contains.

Hadi Hendrawan
Advising CEOs on AI Risk, Authority & Accountability
May 2026

Board Risk Register

1. Hidden Cost Escalation Risk (Reconstruction Cost)

Risk Description
AI systems incur repeated computation due to lack of persistent operational memory, leading to escalating cost per decision.

Root Cause • Context reconstruction per interaction. • Repeated reasoning loops. • Lack of continuity design.	**Exposure** • Margin erosion at scale. • Unpredictable operating cost growth. • Budget misalignment vs actual usage.
Early Indicators • Rising token usage per transaction. • Stable outcomes despite increasing compute. • Difficulty attributing cost to business value.	**Controls** • Define and track **cost per decision**, not just token usage. • Implement **context persistence strategy** (where appropriate). • Set thresholds for **max compute per workflow.** • Audit repeated reasoning loops across systems.

Board Oversight Question

> Do we know how much we are paying to arrive at a single decision and how often it is recomputed?

2. Authority Ambiguity Risk (Unbounded Execution)

Risk Description
AI systems operate without clearly defined authority boundaries, leading to repeated validation cycles and uncontrolled execution behavior.

Root Cause • Undefined decision rights. • Lack of explicit execution limits. • No formal authority architecture.	**Exposure** • Cost inflation via repeated checks. • Inconsistent decision outcomes • Operational unpredictability.
Early Indicators • Frequent re-validation of similar outputs. • Over-reliance on iterative prompting. • Escalation loops without resolution.	**Controls** • Define **authority architecture (residence, boundaries, activation, accountability).** • Enforce **explicit execution thresholds.** • Require **formal instruction triggers** before execution. • Implement override and kill-switch protocols.

Board Oversight Question

Where does authority reside in our AI systems and where does it stop?

3. Governance Cost Multiplication Risk

Risk Description
Governance, compliance, and safety layers increase computational overhead, amplifying cost without proportional value.

Root Cause • Layered validation checks. • Redundant compliance workflows. • Poorly designed control mechanisms.	**Exposure** • Exponential cost growth. • Slower execution without improved outcomes. • Governance becoming a cost center rather than control.
Early Indicators • Significant token usage in validation layers. • Duplicate or overlapping checks. • Increasing latency in decision workflows.	**Controls** • Rationalize governance layers (remove redundancy). • Implement **risk-based control tiers** (not all decisions treated equally). • Define **maximum allowable validation cycles per decision.** • Align governance with **reversibility and risk level.**

Board Oversight Question

> Are our governance controls reducing risk or multiplying cost without increasing control?

4. Decision Quality Dilution Risk

Risk Description
Increased computation and repeated reasoning do not translate into better decisions, leading to inefficiency and potential misjudgment.

Root Cause • Over-reliance on system-generated outputs. • Lack of human decision validation at critical points. • Reprocessing instead of improving reasoning.	**Exposure** • Declining decision quality. • False confidence in system outputs. • Strategic misalignment.
Early Indicators • High compute with minimal improvement in outcomes. • Reduced human challenge or oversight. • Convergence toward a single system-generated option.	**Controls** • Enforce **decision gating mechanisms** (challenge, alternatives, justification). • Track **decision delta vs compute used.** • Require **human validation for high-impact decisions.** • Introduce adversarial review processes.

Board Oversight Question

> Are we improving decisions or just increasing the cost of arriving at them?

5. Continuity & Context Risk (Epistemic Instability)

Risk Description
Lack of persistent memory leads to inconsistent outputs and reliance on reconstructed context.

Root Cause • Stateless model interactions. • Fragmented context storage. • Poor integration between systems.	**Exposure** • Inconsistent decisions across similar scenarios. • Loss of institutional knowledge. • Increased reprocessing cost.
Early Indicators • Different outputs for identical inputs over time. • Frequent need to “re-explain” context. • Dependence on manual context injection.	**Controls** • Implement **structured system-of-record integration.** • Define **context retention policies.** • Standardize prompt/context frameworks. • Audit continuity across workflows.

Board Oversight Question

Are our systems building knowledge or repeatedly reconstructing it?

6. Scaling Misinterpretation Risk

Risk Description
Growth metrics (e.g., token usage, interactions) are interpreted as value creation rather than cost expansion.

Root Cause • Misaligned KPIs (usage vs value). • Lack of cost-to-value mapping. • Vendor-driven metrics.	**Exposure** • Overinvestment in inefficient architectures. • Strategic misdirection. • Inflated perception of ROI.
Early Indicators • Increased usage without revenue or margin improvement. • Focus on activity metrics (tokens, queries). • Lack of outcome-based measurement.	**Controls** • Shift KPIs to **value per decision.** • Link AI usage to **financial outcomes.** • Require ROI justification for scaling initiatives. • Separate **activity metrics vs value metrics.**

Board Oversight Question

Are we measuring activity or value?

7. Accountability & Fiduciary Risk

Risk Description
Responsibility for AI-driven outcomes becomes unclear when systems act autonomously without defined authority structures.

Root Cause • Lack of named accountable owner. • Implicit delegation to systems. • Undefined escalation protocols.	**Exposure** • Legal and regulatory risk. • Governance failure. • Inability to defend decisions under scrutiny.
Early Indicators • Ambiguity in ownership of AI decisions. • Post-incident blame shifting. • Lack of documented authority boundaries.	**Controls** • Assign **named accountable executive for each system.** • Define **authority ownership and escalation paths.** • Document decision rights and limits. • Ensure Board approval for authority expansion.

Board Oversight Question

> Who is accountable when the system acts and is that accountability defensible?

Board Mandate

The Board does not oversee AI execution.

It oversees:

- where authority is granted,
- how it is bounded,
- what it costs when exercised.

Closing Insight

AI introduces a new class of exposure:

Not execution risk.

Design risk.

- Cost is defined before execution.
- Control is defined before execution.
- Accountability is defined before execution.

Autonomy scales. Authority contains.

Hadi Hendrawan
Advising CEOs on AI Risk, Authority & Accountability
May 2026

Board Audit Checklist

Purpose

This checklist is a gating mechanism.

If these conditions are not met:

- authority is being exercised implicitly,
- cost is not governed,
- accountability is not defensible.

1. Authority Definition (Non-Negotiable)

☐ Authority residence is explicitly defined for each AI system.
☐ Execution boundaries are formally documented and enforced.
☐ Activation triggers are structured and auditable.
☐ A named accountable executive is assigned.

If any box is unchecked: authority is undefined.

2. Instruction Formalisation

☐ System execution is triggered only by defined instruction formats.
☐ Conversational inputs cannot trigger operational actions without formalisation.
☐ Instruction → Approval → Execution sequence is preserved where required.

If unchecked: discussion may be executed without authority.

3. Cost Visibility (Decision-Level)

☐ Cost per decision is measured and tracked.
☐ Token usage is mapped to specific workflows.
☐ Repeated computation loops are identified and quantified.

If unchecked: cost is being incurred without visibility.

4. Computation Boundaries

☐ Maximum compute per workflow is defined.
☐ Maximum validation cycles per decision are enforced.
☐ Escalation is triggered when thresholds are exceeded.

If unchecked: systems may expand computation without limit.

5. Governance Efficiency

☐ Governance layers are mapped and non-redundant.
☐ Control intensity is tiered based on risk.
☐ Governance overhead is measured as % of total computation.

If unchecked: governance may be multiplying cost without increasing control.

6. Continuity & Context Integrity

☐ System-of-record integration is defined and enforced.
☐ Context frameworks are standardised across workflows.
☐ Output consistency is tested for identical conditions.

If unchecked: systems are reconstructing context instead of retaining it.

7. Decision Integrity

☐ High-impact decisions require documented alternatives and challenge.
☐ Decision rationale is recorded and auditable.
☐ Human validation is enforced for irreversible or accountability-bearing actions.

If unchecked: decision quality may be degrading without visibility.

8. Scaling Discipline

☐ AI usage metrics are linked to financial outcomes
☐ Activity metrics (tokens, queries) are separated from value metrics
☐ ROI is evaluated at decision level, not system level

If unchecked: scaling may be increasing cost without increasing value.

9. Accountability & Escalation

☐ Every system-executed action has a named accountable owner.
☐ Escalation paths are defined and tested.
☐ Override / kill-switch mechanisms are accessible and enforceable.

If unchecked: accountability is not defensible.

10. Board Control Points

☐ Board has approved authority boundaries for all critical systems.
☐ Authority expansion requires explicit Board approval.
☐ Authority architecture is reviewed periodically.

If unchecked: authority is evolving without oversight.

Audit Outcome Classification

Compliant

All conditions satisfied.
Authority, cost, and accountability are governed.

At Risk

One or more sections partially satisfied.
Implicit authority or cost leakage may exist.

Non-Compliant

Any critical section (1, 2, 3, or 9) fails.

Immediate action required.
System execution may be occurring without defensible authority.

Board Interpretation

This checklist does not assess model performance.

It assesses:

- whether authority is defined before execution,
- whether cost is controlled before scaling,
- whether accountability exists before failure.

Closing Constraint

If these conditions are not met:

The system is not governed.

It is operating on assumed authority.

Autonomy scales. Authority contains.

Hadi Hendrawan
Advising CEOs on AI Risk, Authority & Accountability
May 2026

