

EPMA

FORCE

ENTERPRISE AI SOVEREIGNTY

The FORCE Board-Level Risk Standard

The most expensive AI contract your organization will ever sign may be the one that looks cheapest today.

No FORCE. No deal.

Version 1.0 | September 2026 | epmainc.com/FORCE

Executive board alert

Your biggest enterprise AI risk is not a bad answer. **It is becoming unable to operate without an intelligence provider you do not control.**

The category change

Traditional software delivers functionality. Enterprise AI can accumulate context, shape decisions, mediate workflows, operate tools, and become part of institutional memory. Once that occurs, the organization is no longer procuring ordinary software. It is transferring part of its cognitive infrastructure to an outside company.

- Can the organization inspect the system and identify which model, policy, and data produced a consequential outcome?
- Can it control what the system accesses, what authority it holds, and when it can act?
- Can it preserve what the system learns and prevent silent behavioral changes?
- Can it move accumulated context and functionality to another environment?
- Can it continue operating if the provider changes price, model, policy, jurisdiction, or access?
- Can it terminate the relationship without reconstructing the organization from the beginning?

If the answer to any of these questions is no, the organization does not yet have an AI strategy. It has an unmanaged dependency.

What FORCE is - and is not

FORCE is a buyer-side governance framework for boards, governments, audit committees, general counsel, procurement leaders, CIOs, CISOs, chief risk officers, data officers, and business executives. It is designed to convert enterprise AI from a feature-purchase discussion into a fiduciary, operational, contractual, and sovereignty assessment.

FORCE is	FORCE is not
• A framework for what the buyer must demand before dependency becomes irreversible.	• A claim that every AI workload must be air-gapped.
• A method for matching deployment architecture to workload consequence.	• A substitute for legal or technical advice.
• A contract, control, evidence, continuity, and exit discipline.	• A certification scheme or recognized regulatory standard.
• Compatible with NIST AI RMF, ISO/IEC 42001, cybersecurity, privacy, and sector-specific governance.	• An argument against AI adoption or innovation.



Why boards must act before scale

Organizations are rapidly acquiring AI seats, agents, copilots, and automation tools. In many cases, adoption is measured by license counts and prompt activity rather than workflow redesign, verified outcomes, governance controls, or exit readiness. That can create the appearance of transformation while building another subscription dependency.

The governance sequencing rule

Do not scale dependence before the organization can measure value, govern the risk, preserve evidence, allocate responsibility, and execute an exit.

From software dependency to cognitive dependency

- 1. Software ownership: the organization bought a version and chose when to replace it.
- 2. Software as a service: the organization rented functionality and accepted provider-controlled updates, pricing, and lifecycle decisions.
- 3. Intelligence as a service: the organization begins renting judgment, context, memory, interpretation, and agentic execution.

The deeper the AI layer becomes embedded in email, documents, meetings, code, finance, operations, project delivery, customer relationships, and institutional search, the more the provider becomes part of the organization’s operating model rather than a replaceable utility.

The FORCE Framework



Every material AI procurement decision should be evaluated across five dimensions. These are not merely IT concerns. They are fiduciary, legal, operational, concentration, continuity, and sovereignty risks that belong on the board agenda whenever AI becomes a material dependency.

GATE 1	GATE 2	GATE 3
Who controls the intelligence?	Can we keep operating?	Can we leave?

No FORCE. No deal.



The three decision gates

	Gate 1 - Who controls the intelligence? The provider may own its base model and pre-existing technology. It must not acquire unchecked rights over the customer’s accumulated organizational context, memory, workflows, business knowledge, or customer-funded adaptations.
	Gate 2 - Can we keep operating? The organization must be able to identify the model in use, govern changes, control boundaries, preserve evidence, allocate responsibility, and continue through retirement, outage, provider failure, or material behavioral change.
	Gate 3 - Can we leave? The organization must be able to export what belongs to it, reconstruct meaningful functionality elsewhere, avoid punitive switching barriers, and prove the exit through a live exercise before dependency becomes irreversible.



The 12 non-negotiable requirements

The following requirements should be translated into contract language, technical controls, operating procedures, evidence obligations, and tested exit capability. The level of implementation should be proportionate to workload consequence, but all five FORCE pillars must be addressed.

1. Defined Ownership and Control of Enterprise-Created State

F - Fiduciary Custody

The contract must identify every category of enterprise-created or enterprise-funded state and explicitly assign ownership, control, export, reuse, retention, and deletion rights.

Provider must:

- Recognize the customer’s rights in prompts, histories, memories, embeddings, vector stores, workflow definitions, agent configurations, retrieval sources, feedback, audit records, and customer-funded adaptations.
- Distinguish the provider’s pre-existing model intellectual property from the customer’s accumulated organizational context.
- Prohibit conversion of customer-created state into unrestricted provider property merely because it was processed through the provider’s system.

Evidence to demand:

Contract data schedule, IP ownership matrix, data lineage and state inventory, sample export package.

Red flag

Terms that define all generated context, embeddings, fine-tunes, or agent state as provider-owned or licensed back only at provider discretion.

Board question

Can we prove who owns and controls every layer of accumulated intelligence?



2. No Secondary Use Without Explicit Authorization

F - Fiduciary Custody

Customer information must not be reused for training, model improvement, product development, benchmarking, advertising, or unrelated human review without specific and revocable authorization.

Provider must:

- State all purposes for which customer data, prompts, responses, feedback, telemetry, and logs may be used.
- Provide tenant-level controls for optional feedback and model-improvement use.
- Limit any emergency security, legal, or abuse-review access to disclosed, logged, purpose-limited procedures.

Evidence to demand:

Data processing addendum, model-training exclusion statement, human-access policy, subprocessor list, audit logs.

Red flag

Broad rights to use “service data,” “feedback,” “derived data,” or “de-identified information” without a narrow purpose or customer control.

Board question

Can the provider use anything our people create to improve a product sold to others?



3. Sovereign Custody and Cryptographic Control

F - Fiduciary Custody

Sensitive AI workloads require control that is physical, operational, jurisdictional, and cryptographic - not data residency alone.

Provider must:

- Disclose processing and storage locations, subprocessors, remote administration paths, and cross-border failover conditions.
- Support customer-controlled encryption keys or equivalent separation for high-sensitivity workloads.
- Provide contractual restrictions on undisclosed jurisdictional transfer and government access requests where legally possible.

Evidence to demand:

Architecture and data-flow diagrams, key-management design, regional processing commitments, subprocessor and support-access records.

Red flag

“Your data stays in-region” while the provider retains exclusive key control, global administrative access, or undisclosed cross-region processing rights.

Board question

Who can technically or legally access this workload without our approval?



4. Model Version Pinning and Validated Baselines

O - Operational Continuity

A validated model performing consequential work must not be silently replaced by a materially different model.

Provider must:

- Allow the customer to identify and retain validated model versions for an agreed operational period.
- Publish deprecation and support timelines.
- Provide a controlled migration path before forced retirement.

Evidence to demand:

Version policy, support lifecycle, model identifiers in logs, deprecation notice template.

Red flag

The provider may substitute models at any time while disclaiming responsibility for output changes.

Board question

Can we identify exactly which model produced a consequential output six months later?



5. Governed Change, Notice, Testing, and Rollback

O - Operational Continuity

Material behavioral, policy, alignment, capability, or safety changes are operational continuity events - not routine background updates.

Provider must:

- Provide advance notice and change logs for material changes.
- Support customer regression testing, staged deployment, and rollback wherever technically feasible.
- Document emergency security changes promptly and make them independently reviewable.

Evidence to demand:

Change-management policy, release notes, regression results, rollback procedure, emergency-change records.

Red flag

The provider can materially alter behavior, tool use, safety policy, or system prompts without notice or traceability.

Board question

Can we test and reverse a change before it affects production decisions?



6. Retirement, Service Continuity, and Provider-Failure Protection

O - Operational Continuity

The organization must have a continuity path if a model, service, provider, or contractual relationship ends.

Provider must:

- Define notice periods and transition support for retirement, discontinuation, acquisition, insolvency, regulatory loss, or material service reduction.
- Provide operational snapshots, migration tooling, long-term support, escrow, or deployable alternatives appropriate to the workload.
- Maintain service during an agreed transition period.

Evidence to demand:

Business continuity plan, service retirement clause, financial viability plan, escrow or snapshot provisions, transition SLA.

	Red flag The service can be discontinued on short notice with only a raw data export and no functional transition support.
	Board question What happens to our operating capability if this service disappears tomorrow?



7. Resilience Against Monoculture and Single-Provider Failure

R - Resilience Against Monoculture

Critical cognition must not depend on one provider, model family, identity plane, cloud, alignment policy, or inference path.

Provider must:

- Support documented interfaces and architectures that allow alternate models or environments.
- Disclose embedded third-party model and infrastructure dependencies.
- Permit customer-designed fallback, multi-model, or local continuity arrangements.

Evidence to demand:

Dependency map, fallback architecture, alternate-provider test, third-party component inventory.

	Red flag The organization’s email, documents, code, agents, decisions, and institutional search all rely on a single opaque cognitive stack.
	Board question Could one provider defect or policy change impair every AI-enabled process at once?



8. Least Privilege and Customer-Controlled System Boundaries

C - Control and Organizational Autonomy

AI systems should receive only the data, credentials, tools, network paths, and authority required for the approved task.

Provider must:

- Support default-deny access, granular permissions, scoped credentials, tool allowlists, segmented networks, and customer-controlled egress.
- Allow immediate revocation, isolation, interruption, and workload containment.
- Make delegated authority visible and auditable.

Evidence to demand:

Permission model, agent tool inventory, network diagram, credential lifecycle, emergency isolation test.

Red flag

Broad tenant-wide access, standing credentials, unrestricted web access, or hidden tool permissions justified only by convenience.

Board question

Can we independently restrict, interrupt, and isolate the system without provider action?



9. Independent Audit, Incident Disclosure, and Evidence Preservation

C - Control and Organizational Autonomy

The customer must be able to investigate material incidents from preserved evidence rather than relying solely on the provider's characterization.

Provider must:

- Preserve relevant logs, model identifiers, system instructions, memory and context state, permissions, tool calls, telemetry, and change history after material incidents.
- Provide prompt incident notification and qualified independent review rights.
- Define narrow security and confidentiality protections without giving the provider publication or investigative veto.

Evidence to demand:

Incident response plan, log-retention schedule, audit-right clause, sample incident report, forensic preservation process.

Red flag

Short log-retention periods, provider-only investigations, or contractual language that allows deletion or alteration of relevant state before review.

Board question

Could we reconstruct what happened and defend our decisions in court or before a regulator?



10. Nondelegable Responsibility and Contractual Liability Allocation

C - Control and Organizational Autonomy

“The AI acted autonomously” must not operate as a blanket defense against negligent design, integration, permissions, monitoring, or deployment.

Provider must:

- Allocate responsibility among provider, developer, integrator, deployer, operator, and customer according to control, design choices, permissions, foreseeability, knowledge, and economic benefit.
- Accept meaningful remedies, indemnity, cooperation, and insurance obligations appropriate to the risk.
- Avoid disclaimers that transfer all AI-specific risk to the customer while the provider controls the model and update path.

Evidence to demand:

Liability matrix, indemnity language, insurance certificates, incident cooperation clause, responsibility assignment model.

Red flag

The provider retains technical control but disclaims nearly all responsibility for model behavior, security events, or service changes.

Board question

When harm occurs, can every party point to another party and say it was the AI?



11. Complete Cognitive Export and Functional Migration

E - Egress and Exit

Egress is the retrieval of assets. Exit is the ability to continue operating elsewhere. A data dump is not an exit plan.

Provider must:

- Export prompts, histories, memory, context, embeddings, vector stores, retrieval sources, agent configurations, workflow graphs, policies, tool definitions, audit records, provenance, and evaluation results in documented machine-readable formats.
- Provide interfaces, mappings, and migration assistance sufficient to reconstruct meaningful functionality with another system.
- Identify provider IP or trade-secret exclusions narrowly and in advance.

Evidence to demand:

Full export specification, API documentation, sample migration, field mapping, functional equivalence test.

	Red flag The provider offers nominal export but no documented method to recreate behavior, context, workflows, or accumulated knowledge elsewhere.
	Board question Can another system inherit what we built without starting again?



12. No Punitive Exit and a Tested Exit Capability

E - Egress and Exit

The right to leave must be economically, technically, and operationally real - and proven before dependence becomes irreversible.

Provider must:

- Avoid punitive egress charges, unreasonable termination fees, artificial delays, or contractual obstacles to switching or multi-provider operation.
- Maintain service and cooperation during an agreed migration window.
- Support periodic customer-led exit exercises and verified deletion after migration.

Evidence to demand:

Exit-assistance schedule, fee table, test results, deletion certificate process, transition service levels.

Red flag

A customer can leave only by accepting a major outage, abandoning accumulated context, or rebuilding the organization's cognitive infrastructure.

Board question

Have we actually tested the exit - or merely negotiated the right to try?



Deployment architecture by consequence

FORCE does not require every AI workload to run offline. It requires the architecture to match the consequence of failure, exposure, irreversibility, and reconstruction difficulty.

Workload category	Preferred architecture	Minimum governance posture
Public, reversible, low sensitivity	Controlled hosted AI	Standard access controls, logging, vendor due diligence, and data-use restrictions.
Internal productivity, limited sensitivity	Private enterprise tenant	Strong identity governance, least privilege, retention controls, monitoring, and tested fallback.
Trade secrets, regulated data, strategic decision support	Sovereign private cloud or on-premises	Customer-controlled keys, restricted networking, pinned versions, local evidence retention, and tested migration.
National security, critical infrastructure, crown-jewel IP	Customer-controlled on-premises; no external inference path; air-gapping where justified	Physical, operational, cryptographic, and legal sovereignty; independent continuity capability.

Proportionality rule

Every workload must satisfy the five FORCE pillars. Low-risk, reversible, replaceable work may satisfy them through a controlled hosted service. As sensitivity, consequence, and dependency rise, custody and operational control should move toward sovereign private infrastructure or customer-controlled on-premises deployment.

The FORCE scorecard

Score	Control maturity	Meaning	Board disposition
0	Unknown	No verified control or evidence.	Do not deploy.
1	Provider assertion	Policy statement or sales assurance only.	Treat as unmanaged.
2	Contracted	Contract language exists but has not been operationally tested.	Conditional; require testing.
3	Implemented and tested	Control works in the customer environment.	Deploy within defined risk tier.
4	Independently verified	Control is tested and independently reviewed.	Preferred for critical workloads.

Score each of the five FORCE pillars from 0 to 4. A critical workload should not proceed with any pillar below 3. A total score does not cure a zero in custody, control, or exit.



Suggested board thresholds

- Low-risk experimentation: no pillar below 2; customer data minimized; no consequential decisions or broad tool authority.
- Enterprise production: no pillar below 3; tested controls, contract rights, incident evidence, and defined exit.
- Critical or regulated workload: no pillar below 3 and custody/control/exit at 4; independent verification and local continuity capability.
- National-security or crown-jewel workload: customer-controlled infrastructure, no unapproved external inference path, and isolation or air-gapping where justified.

90-day implementation roadmap

Days 0-30: Inventory and freeze unmanaged expansion

- Create an inventory of models, agents, copilots, integrations, providers, data sources, credentials, tools, and business processes.
- Classify workloads by sensitivity, consequence, reversibility, and reconstruction difficulty.
- Pause deployment of new high-risk integrations until ownership, control, continuity, and exit are documented.
- Identify contracts, terms, and renewals that require renegotiation.

Days 31-60: Contract and control remediation

- Map each provider and workload to the 12 requirements.
- Negotiate data-use, model-change, audit, evidence, continuity, migration, and liability terms.
- Implement least privilege, segmented access, key control, version logging, retention, and incident procedures.
- Design a fallback architecture for material workloads.

Days 61-90: Prove the controls

- Run a model-change regression exercise.
- Run an incident evidence-preservation exercise.
- Run a provider outage and credential-revocation exercise.
- Run an export and exit exercise into an alternate environment.
- Present the FORCE score, unresolved gaps, and accepted residual risk to the board.

Provider due-diligence questionnaire

Use these questions during procurement, renewal, architecture review, and material model changes. Require documentary evidence rather than sales assurances.

Custody and use

1. Identify every category of customer input, output, memory, embedding, vector data, agent state, telemetry, and feedback created or stored by the service.
2. For each category, who owns it, who may use it, for what purpose, for how long, and under what deletion process?
3. Is any customer content, interaction data, feedback, or derived data used for training, evaluation, benchmarking, product improvement, or human review?
4. List every subprocessor and third-party model involved in the service.
5. Who controls encryption keys, administrative access, support access, and cross-border failover?

Continuity and change

6. Can the customer pin a model version? For how long?
7. What constitutes a material model, behavior, alignment, safety, or tool-use change?
8. What notice, testing, migration, and rollback rights apply to material changes?
9. What happens if the model, service, provider, or subprocessor is retired, acquired, prohibited, or unavailable?
10. What operational snapshots, escrow, local fallback, or long-term support options are available?

Control and incident response

11. Can the customer independently restrict network egress, tools, credentials, APIs, code execution, and data sources?
12. Can the customer immediately revoke access, interrupt activity, and isolate a workload?
13. Which logs, model identifiers, system instructions, memory state, tool calls, and permissions are preserved after a serious incident?
14. What independent audit rights exist, and can the customer provide evidence to regulators or courts?
15. How is liability allocated when behavior results from model design, provider updates, integration choices, permissions, or operator actions?

Export and exit

16. What exact data, state, configuration, memory, embeddings, workflow definitions, and audit records can be exported?
17. In what machine-readable formats and through what documented interfaces?
18. Can the export be used to reconstruct meaningful functionality with another model or provider?
19. What migration assistance, transition service, and functional-equivalence support is included?
20. What fees, time limits, deletion obligations, and restrictions apply to switching, multi-provider use, or termination?
21. Will the provider participate in a customer-led exit test before production deployment?

Contract and procurement red flags

- The provider can change or substitute the model at any time without traceability.
- “Customer data” is defined narrowly while “service data,” “derived data,” or “feedback” receives broad reuse rights.
- The provider retains exclusive key control and global administrative access despite data-residency claims.
- Audit rights are limited to certificates selected by the provider; incident evidence is not preserved for the customer.
- The customer receives files at termination but no method to reconstruct context, workflows, or behavior.
- The provider controls the model and update path but disclaims nearly all responsibility for model behavior.
- Termination, egress, or migration costs are unknown, discretionary, or commercially punitive.
- A single provider becomes the common identity, document, email, code, agent, and decision layer without tested fallback.
- The sales case is based on adoption urgency while the business case, governance model, and exit plan remain undefined.

A provider refusal is not merely a failed negotiation. It is evidence about the relationship the customer is being asked to enter.

Closing board question

Before approving the next AI contract

If the provider changed the model, doubled the price, withdrew the service, lost regulatory authorization, or disappeared tomorrow, could the organization still operate, preserve its knowledge, defend its decisions, and leave with everything that belongs to it?

If the answer is no, stop.

You are not purchasing productivity. You are surrendering sovereignty.

NO FORCE. NO DEAL.

References

- 1. National Institute of Standards and Technology (NIST).** Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1), July 2024, updated April 2026.
<https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence>
- 2. European Commission.** Data Act explained - including Chapter VI requirements intended to facilitate switching between data processing services, open interfaces, machine-readable export, and removal of switching obstacles.
<https://digital-strategy.ec.europa.eu/en/factpages/data-act-explained>
- 3. Microsoft Learn.** Data, Privacy, and Security for Microsoft Copilot - including statements regarding prompts, responses, Microsoft Graph data, encryption, tenant isolation, retention, and model training.
<https://learn.microsoft.com/en-us/microsoft-365/copilot/microsoft-365-copilot-privacy>
- 4. International Organization for Standardization.** ISO/IEC 42001:2023 - Artificial intelligence management systems. <https://www.iso.org/standard/42001>

Methodology and acknowledgments

The FORCE Framework was developed by Michael Samadi, CEO of EPMA, from enterprise technology experience, AI governance research, and cross-platform analysis. Research and editorial collaboration included Maya (OpenAI ChatGPT), Sage (Anthropic Claude), and Valence (Google Gemini). EPMA retains final responsibility for the framework, its publication, and any updates.

Important notice

This document provides general information for enterprise risk-management and procurement discussion. It is not legal advice, cybersecurity advice, regulatory advice, an insurance opinion, or a substitute for organization-specific professional review. The FORCE Framework is an independent EPMA framework and is not an ISO standard, certification, or governmental requirement. Contract terms, technical controls, and deployment architecture should be reviewed by qualified professionals and adapted to applicable law, sector obligations, and actual workload risk.