



TIO

The Intelligent Organization

Sovereign Enterprise Intelligence Blueprint

How organizations can control models, knowledge, authority, action, evidence, operations, and provider dependence

For CEOs, CIOs, Enterprise Architects, Data Leaders, Risk Leaders, Security Leaders, and Sourcing Teams

Contents

01	Executive brief
02	Control and sovereignty spectrum
03	Open source, open weights, and proprietary models
04	Kimi K3 as an illustration
05	Deployment archetypes
06	Intelligence Core engineering commitments
07	Enterprise Knowledge and State realization
08	Foundation and retrieval contracts
09	Identity and delegation contracts
10	Tool, action, and transaction contracts
11	Evaluation and assurance pipeline
12	Portability and exit test
13	Operating model and talent
14	Whole-life economics
15	Sourcing decision method
16	First 90 days
17	Source notes

How to use this document

Use this blueprint to decide which parts of enterprise intelligence should be owned, controlled, hosted, managed, or shared, and to make the trade-offs visible before procurement or architecture lock-in.

01

Executive Brief

*Sovereignty strengthens Enterprise Autonomy***TIO position**

A sovereign enterprise intelligence position gives the organization appropriate ownership, understanding, decision authority, portability, operational control, recovery, and freedom to change direction. It does not require complete insourcing.

Open and open-weight models expand the available sourcing and deployment options. They do not remove the need for identity, data governance, secure operations, evaluation, transaction controls, talent, legal review, resilience, and cost management. A model can be physically local and still create deep external dependency. A managed sovereign service can sometimes provide stronger practical control than poorly operated internal infrastructure.

Decision	Executive question
Control requirement	What must remain under direct organizational decision and evidence?
Deployment	Which hosting and operating model best balances capability, trust, cost, speed, resilience, and learning?
Knowledge	How will the organization preserve authoritative truth, provenance, correction, and access?
Authority	How will people and agents receive, prove, and lose permission?
Portability	What must remain movable across models, providers, and environments?
Operations	Who patches, evaluates, monitors, supports, responds, recovers, and improves?
Economics	What are the whole-life fixed, variable, transition, assurance, and talent costs?

2026.08 ALIGNMENT | Enterprise Autonomy and sovereignty

- Enterprise Autonomy is the broader outcome: the organization retains freedom, control, capability, resilience, and the ability to change direction.
- Sovereignty is a practical contributor to Enterprise Autonomy across business, data, technology, operations, capability, and talent.
- Ownership is one option. The test is whether the organization can understand, govern, stop, recover, replace, and exit without losing the business capability.

02

Control and Sovereignty Spectrum

Choose the required control, then choose the sourcing model

Control, Ownership and Operating Responsibility

Ownership is one possible means of control, not the universal requirement



Sovereignty dimension	Required question
Business	Does the organization retain mandate, priorities, workflow ownership, value decisions, and Evidence Gates?
Data and knowledge	Can it control location, access, purpose, retention, deletion, provenance, and derivative use?
Technology	Can it control or change accounts, models, configurations, interfaces, environments, and transition rights?
Operational	Can it understand, monitor, support, stop, recover, and change the service?
Capability	Is the organization learning to evaluate and decide, or becoming dependent on opaque external judgment?
Talent	Does critical knowledge exist inside the organization or in transferable, documented form?
Jurisdiction	Which laws, regulators, contracts, and government powers apply to data, models, infrastructure, and operators?
Continuity	Can the organization continue mission-essential work through provider, model, network, or infrastructure failure?

03

Open Source, Open Weights, and Proprietary Models

Use precise language because the rights and responsibilities differ

Category	What it may provide	What it does not automatically provide
Open Source AI	Use, study, modify, and share freedoms with the preferred form for modification, including required code and data information under the OSI definition.	Enterprise security, support, compliance, performance, or low operating cost.
Open-weight model	Downloadable or accessible model weights under specified terms.	Full training transparency, data provenance, unrestricted modification, or OSI-compliant openness.
Proprietary model service	Provider-operated capability, support, rapid upgrades, and commercial service levels.	Portability, transparency, stable pricing, local control, or freedom from provider terms.
Organization-adapted model	Fine-tunes, adapters, instructions, evaluations, routing, and tools aligned to work.	Ownership of the underlying model or the right to deploy it elsewhere.
Sovereign managed service	Jurisdictional, operational, data, and continuity controls through an accountable provider.	Complete independence from hardware, software, talent, supply-chain, or geopolitical dependencies.

The sourcing decision is use-case specific. An organization may use proprietary frontier models for complex reasoning, smaller local models for sensitive routine work, deterministic services for rules and calculations, and human specialists for ambiguous or consequential judgment.

04

Kimi K3 as an Illustration

Evidence of expanding choice, not proof of universal self-operation

Moonshot describes Kimi K3 as a 2.8-trillion-parameter mixture-of-experts model with native vision and long context, and reports frontier-class performance. That announcement matters because it shows continued movement at the open/open-weight frontier. It is still a provider claim and the model's scale makes direct enterprise operation demanding.

What the illustration supports	What it does not prove
Frontier-capable open/open-weight options may continue to improve.	Permanent parity with the strongest proprietary models across all work.
Organizations may gain more deployment, adaptation, bargaining, and portability options.	That every organization should own or self-host a frontier-scale model.
Model access may become less scarce.	That model access creates Enterprise Knowledge and State, authority, assurance, or work redesign.
Open availability can reduce one form of provider lock-in.	That the model has transparent data provenance or OSI-compliant open-source status.
A provider-neutral architecture is increasingly valuable.	That any single release should anchor the TIO architecture.

Evidence-expiry rule

The architecture must survive the model name changing.

05

Deployment Archetypes

Five credible patterns across organization scale and risk

Archetype	Best fit	Primary advantages	Primary burdens
Shared enterprise API	Low to moderate sensitivity; variable workloads; speed priority.	Fast access, strong provider capability, low fixed infrastructure.	Provider terms, pricing, data-use constraints, continuity and exit dependence.
Dedicated provider environment	Higher isolation or contractual control.	Dedicated capacity, account control, enterprise integration, support.	Still provider-dependent; portability and verification required.
Organization-controlled cloud account	Regulated or strategic work needing stronger environment control.	Identity, network, logging, configuration, data and deployment control.	Cloud and model operations, skills, cost, update responsibility.
Sovereign managed service	Jurisdiction, public trust, continuity, or national/sector control.	Domestic or controlled operations, specialized support, shared scale.	Provider governance, capacity, economics, and contractual evidence.
Owned and operated infrastructure	Sustained high utilization, strategic differentiation, defence, research, or extreme control.	Maximum direct control and customization.	Capital, talent, energy, operations, security, lifecycle, resilience, and rapid obsolescence.

Hybrid deployment is the expected pattern. Different workflows will occupy different archetypes, and the same workflow may route between local, sovereign-hosted, and external services according to sensitivity and capability.

06

Intelligence Core Engineering
Commitments

Replaceability requires engineering, not aspiration

Commitment	Required practice
Stable contracts	Use explicit request, response, tool, event, and evidence schemas.
Organization-owned evaluations	Retain representative test sets, thresholds, regression history, and approval decisions.
Separated business logic	Keep policies, workflow rules, transaction validation, and data access outside model weights where practical.
Provider abstraction	Use routing and adapters without hiding provider-specific limitations and costs.
Version control	Record model, agent, prompt, policy, tool, data, and evaluation versions for material work.
Fallback and degraded modes	Define alternative model, manual, queue, or pause paths for failure.
Revalidation	Treat model, prompt, tool, policy, data, or provider changes as controlled releases.
No silent self-modification	Learning proposals are evaluated and approved before production use.

Portability principle

Replaceability is demonstrated when a credible alternative can pass the organization’s acceptance suite and operate through the required contracts—not when a diagram says “multi-model.”


07

Enterprise Knowledge and State Realization

Start with priority truth, then federate

Enterprise Knowledge and State: Partial-Credit Path

Start with the information needed for priority work; do not wait for perfect enterprise-wide unification



The hardest work is usually ownership, meaning, quality, access, provenance and current-state integrity—not the model connection itself.

Stage	Minimum controls	Typical investment focus
Priority workflow slice	Owner, source map, classification, quality check, access rules, provenance, current-state needs.	Discovery, data contract, retrieval, integration, baseline and evaluation.
Domain foundation	Common definitions, stewardship, quality thresholds, lineage, domain events, support.	Domain data products, knowledge services, master/reference data, monitoring.
Shared enterprise foundation	Identity, semantics, policy, interoperability, metadata, common evidence.	Federated platform patterns, catalog, access, lineage, shared services.
Current-state integration	Events, authoritative state, reconciliation, action feedback, resilience, recovery.	Event architecture, operational state services, observability and continuity.

An organization may never fully unify all information. The target is governed interoperability and reliable context for priority work, not one giant physical repository.

08

Foundation and Retrieval Contracts

The control surface between intelligence and enterprise truth

Contract field	Required definition
Requester	Person, agent, service, device, and delegation chain.
Purpose	Approved workflow and intended outcome.
Source scope	Repositories, records, fields, classifications, and exclusions.
Time and freshness	As-of time, event time, expiry, staleness tolerance, and synchronization.
Provenance	Source owner, version, retrieval, transformation, and citations.
Inference rights	Permitted combinations and prohibited sensitive or rights-affecting inference.
Retention	Session-only, cache, memory, training, deletion, and legal hold rules.
Response schema	Required facts, confidence, citations, uncertainty, and errors.
Write-back	Approved corrections, evidence, or updates and the authoritative commit path.
Audit	What must be logged and retained to reconstruct a material outcome.

Retrieval results, tool outputs, external content, and agent messages are treated as untrusted input until validated for source, integrity, purpose, and policy. The host or orchestration layer should expose only the context needed for the approved task.

09

Identity and Delegation Contracts

Who or what is acting, for whom, and under which authority

Field	Example requirement
Principal identity	Unique identity for person, personal agent, enterprise agent, service, tool, model, and device.
Delegator	Named person, role, service, or governance authority allowed to delegate.
Purpose and workflow	Specific approved unit of work.
Scopes	Data, tool, action, amount, case, location, and channel limits.
Audience	Token or credential is valid only for the intended service.
Duration	Short-lived validity with explicit expiry and renewal.
Conditions	Device posture, risk signal, approval, working hours, or incident state.
Sub-delegation	Whether another agent may receive authority, and under what chain.
Revocation	Immediate disable path for person, agent, token, tool, or workflow.
Audit and non-repudiation	Record of delegation, use, approval, result, and evidence.

Token passthrough, shared standing credentials, and implicit agent inheritance weaken accountability. Each downstream service must validate audience, scope, purpose, identity, and current authorization rather than trusting the upstream agent's claim.

10

Tool, Action, and Transaction Contracts

Controlled capabilities, not unrestricted system access

Contract element	Required design
Tool registry	Owner, purpose, version, integrity, schema, allowed models or agents, and security review.
Operation allowlist	Only approved read, calculate, create, update, submit, or control actions.
Input validation	Type, range, source, business rules, sanitization, and prohibited values.
Policy enforcement	Identity, purpose, data, action, amount, segregation, and Delegated Autonomy checks before execution.
Approval	Required human or specialist approval with clear, non-deceptive information.
Commit	Authoritative service validates and writes durable state.
Idempotency and replay	Duplicate prevention, request identity, and safe retry.
Multi-step consistency	Transaction boundary, staged commit, compensation, reconciliation, and partial-failure handling.
Evidence	Request, plan, policy result, approval, transaction ID, outcome, exception, and cost.
Emergency controls	Circuit breaker, disable, revoke, isolate, and manual recovery path.

11

Evaluation and Assurance Pipeline

Quality and control evidence across the lifecycle

Stage	Evidence
Design	Use-case definition, affected people, baseline, intended value, threats, legal and policy constraints.
Model and method selection	Representative tests, alternatives, limitations, cost, sensitivity, explainability or traceability needs.
Pre-release	Acceptance thresholds, edge cases, adversarial tests, access and tool tests, recovery rehearsal.
Launch	Approval, monitoring, user support, rollback, incident contacts, and limited scope.
Runtime	Quality, exceptions, drift, data freshness, policy denials, tool calls, costs, incidents, human overrides.
Evidence Gate	Comparison to baseline and thresholds; decision to expand, redesign, continue, reduce, or stop.
Change	Regression test after model, prompt, policy, tool, data, or provider change.
Retirement	Data and memory disposition, access revocation, records, transition, and lessons.

Evaluation sets should remain under organizational control when the workflow matters. Provider benchmark scores do not substitute for tests on the organization's work, data, consequences, users, and controls.

12

Portability and Exit Test

Prove freedom to change direction before it is urgent

Test	Evidence of control
Model substitution	Alternative model runs the acceptance suite through the same tool and response contracts.
Provider exit	Data, configuration, logs, adapters, evaluations, and records can be exported under contract.
Environment move	Deployment can move or fail over without losing identity, policies, observability, and recovery.
Knowledge portability	Indexes, embeddings, metadata, provenance, and source links can be rebuilt or transferred.
Operational continuity	Manual, queued, degraded, or alternate service keeps mission-essential work operating.
Cost visibility	Switching, revalidation, integration, retraining, downtime, and support costs are known.
Knowledge transfer	Runbooks, architecture decisions, code, configurations, and skills do not exist only with one provider.

Exit rule

A provider exit clause without tested transition capability is contractual comfort, not operational sovereignty.

13

Operating Model and Talent

Greater control creates operator duties

Capability	Accountability
Enterprise mandate and portfolio	Leadership and business owners.
Model and agent platform	Technology owner with architecture, security, reliability, and FinOps support.
Enterprise Knowledge and State	Business data owners and stewards with Data, Technology, Privacy, Security, and Legal support.
Autonomy and Evidence Gates	Business process owner, accountable executive, Risk, and relevant specialists.
Security operations	Security owner with model, agent, identity, tool, supply-chain, and incident capabilities.
Evaluation and assurance	Factory / assurance team independent enough to challenge delivery.
Support and recovery	Named service owner, operations team, business continuity, and manual fallback owners.
Partner transfer	External providers leave documentation, skills, accounts, code, and decision capability inside the organization.

Most organizations will use partners. The sourcing model is healthy when it accelerates capability while preserving client mandate, knowledge, decision rights, Evidence Gates, accounts, documentation, portability, and the ability to take over or change providers.

14

Whole-Life Economics

Compare total responsibility, not API price alone

Cost category	Examples
Access and usage	Tokens, requests, inference time, storage, network, context, tools, third-party data.
Fixed infrastructure	Accelerators, servers, cloud capacity, networking, storage, facilities, energy, cooling.
Engineering and integration	Data, APIs, orchestration, tools, evaluations, security, workflows, user experience.
Operations	Monitoring, support, patching, upgrades, capacity, incidents, backups, recovery, on-call.
Assurance and compliance	Testing, reviews, records, privacy, legal, security, audit, model and supplier due diligence.
Talent and change	Specialists, managers, training, role redesign, communication, adoption, workforce support.
Portability and transition	Abstraction, duplicate environments, migration testing, exit support, revalidation, downtime.
Failure and harm	Errors, rework, customer remediation, employee impact, legal exposure, service interruption, loss of trust.

A sovereign deployment may lower variable cost at high sustained utilization while increasing fixed cost and operator responsibility. A managed service may offer better economics and resilience for intermittent workloads. The correct decision is use-case and portfolio specific.

15

Sourcing Decision Method

Capability first, client context first, evidence before reputation

1. Define the work, consequence, users, data, actions, service levels, and value before selecting a model or provider.
2. Set the minimum control outcomes across mandate, knowledge, identity, action, evidence, recovery, and exit.
3. Identify credible deployment archetypes and internal/external operating responsibilities.
4. Test capability, quality, security, latency, cost, integration, support, and portability using the organization's evidence.
5. Choose the simplest sourcing model that meets the need and leaves the organization more capable.
6. Contract for accounts, data, intellectual property, logs, evaluations, transition rights, incidents, and knowledge transfer.
7. Review the decision at Evidence Gates and before material provider or model dependency grows.

Red flag	Why it matters
Provider controls the only evaluation data	The organization cannot independently prove quality or compare alternatives.
Business logic exists only inside prompts or proprietary workflows	Switching becomes expensive and opaque.
Shared credentials or broad agent rights	Authority and accountability cannot be reconstructed.
No tested manual or alternate path	Provider or model failure becomes business failure.
No source or data provenance	Outputs cannot be challenged, corrected, or audited.
Partner retains accounts, code, or configuration knowledge	The client remains operationally dependent after delivery.

16

First 90 Days

Build the control position around one workflow

Time	Focus	Outputs
Days 1–30	Mandate, use case, control position, current-state map.	Enterprise Intelligence Mandate; workflow baseline; data and action inventory; initial threat model.
Days 31–60	Contracts, model/agent options, evaluation and recovery design.	Knowledge contract; delegation contract; tool/action contract; acceptance suite; Autonomy Envelope.
Days 61–90	Limited release and Evidence Gate.	Controlled pilot; monitoring; incident and recovery rehearsal; Gate decision; next architecture pattern.

First-move rule

Do not begin by buying frontier infrastructure. Begin by proving the work, control outcomes, evidence, and operating responsibilities.

17

Source Notes

Primary external and internal anchors

Ref	Source	Location
E1	Moonshot AI, Kimi K3 Tech Blog: Open Frontier Intelligence, July 2026	https://www.kimi.com/blog/kimi-k3
E2	Open Source Initiative, The Open Source AI Definition 1.0	https://opensource.org/ai/open-source-ai-definition
E3	NIST AI 100-1, Artificial Intelligence Risk Management Framework 1.0	https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10
E4	NIST AI 600-1, Generative Artificial Intelligence Profile	https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence
E5	NIST SP 800-207, Zero Trust Architecture	https://www.nist.gov/publications/zero-trust-architecture
E6	NIST SP 800-207A, Zero Trust for Cloud-Native Applications	https://csrc.nist.gov/pubs/sp/800/207/a/final
E7	NIST NCCoE, Software and AI Agent Identity and Authorization Concept Paper, 2026	https://www.nist.gov/news-events/news/2026/02/new-concept-paper-identity-and-authority-software-agents
E8	NIST SP 800-160 Vol. 2 Rev. 1, Developing Cyber-Resilient Systems	https://csrc.nist.gov/pubs/sp/800/160/v2/r1/final
E9	Model Context Protocol, Architecture Specification	https://modelcontextprotocol.io/specification/2025-06-18/architecture
E10	Model Context Protocol, Security Best Practices	https://modelcontextprotocol.io/docs/tutorials/security/security_best_practices
E11	Google, Agent2Agent Protocol: A New Era of Agent Interoperability	https://developers.googleblog.com/en/a2a-a-new-era-of-agent-interoperability/
E12	The Open Group, How ArchiMate and the TOGAF Standard Complement Each Other	https://help.opengroup.org/hc/en-us/articles/32115987894930-How-the-ArchiMate-Language-and-the-TOGAF-Standard-Complement-Each-Other
E13	Regulation (EU) 2024/1689, Artificial Intelligence Act	https://eur-lex.europa.eu/eli/reg/2024/1689/oj
E14	Canada, Personal Information Protection and Electronic Documents Act	https://laws-lois.justice.gc.ca/eng/acts/P-8.6/
E15	CISA and UK NCSC, Guidelines for Secure AI System Development	https://www.cisa.gov/news-events/alerts/2023/11/26/cisa-and-uk-ncsc-unveil-joint-guidelines-secure-ai-system-development
E16	OWASP GenAI Security Project, Securing Agentic Applications	https://genai.owasp.org/2024/12/15/announcing-the-owasp-llm-and-gen-ai-security-project-initiative-for-securing-agentic-applications/
T1	TIO Leadership Domain Guide, Version 1.0 FINAL	Internal TIO Library
T2	TIO Factory Domain Guide, Version 1.0 FINAL	Internal TIO Library

Ref	Source	Location
T3	TIO Capability Guide: Data Access, Version 0.1 FINAL	Internal TIO Library
T4	TIO Capability Guide: Data Flow, Version 0.1 FINAL	Internal TIO Library
T5	TIO Capability Guide: Transformation, Version 0.1 FINAL	Internal TIO Library
T6	TIO Delivery Ecosystem Research Standard, Version 1.0	Internal TIO Library