



TIO
The Intelligent Organization

Enterprise Intelligence Mandate

A practical executive charter for Alpha, its control position, authority, evidence, and first operating boundaries

For CEOs, owners, executive teams, accountable sponsors, CIOs, and governance leaders

Contents

01	How to use this mandate
02	Mandate statement
03	Outcomes and control position
04	Scope and non-goals
05	Governance and accountability
06	Human authority
07	Enterprise Knowledge and State
08	Models, agents, tools, and sourcing
09	Autonomy posture
10	Action and transaction
11	Assurance, incidents, and recovery
12	People and Workforce Charter
13	First 90 days
14	Decision and approval record
15	Source notes

How to use this document

Complete this mandate before broad platform decisions or high-consequence agent deployments. Adapt every bracketed or illustrative item to the organization and obtain the required executive, legal, risk, privacy, security, and workforce input.

01

How to Use This Mandate

Set the enterprise boundaries before the architecture grows

The Enterprise Intelligence Mandate is the leadership charter for Alpha. It defines why the capability exists, what the organization must control, what work is in scope, which decisions remain human, how Delegated Autonomy is granted, how evidence is reviewed, and who is accountable.

Alpha combines an Intelligence Core with governed Enterprise Knowledge and State. The Intelligence Core supplies replaceable models, agents, methods, tools, evaluations, and human expertise. Enterprise Knowledge and State supplies governed context, intellectual property, authoritative truth, evidence, meaning, and current operating condition.

Boundary

This mandate does not authorize a specific model, provider, data source, agent, or autonomous workflow. Those decisions require the relevant design, risk, data, security, legal, procurement, and Evidence Gate approvals.

2026.08 ALIGNMENT | Mandate outcomes and decisions

- The mandate should state the intended Intelligence, Automation, and Enterprise Autonomy outcomes.
- It should also define the transition route for material units of work and the Assurance and Operations responsibilities after launch.
- Any Delegated Autonomy remains separate, bounded, evidence-gated, reviewable, revocable, and recoverable.

02

Mandate Statement

Plain-English executive direction

Proposed mandate

Our organization will build and govern Alpha: a coherent enterprise intelligence capability that helps people and systems understand, decide, coordinate, act, learn, and adapt. Alpha will use governed Enterprise Knowledge and State, approved intelligence, explicit identity and authority, controlled transaction services, continuous evidence, and named human accountability. We will expand machine authority only when the evidence justifies it.

Alpha is the TIO reference name. The organization may choose another operating name. The mandate applies to the complete system rather than to one model or platform.

03

Outcomes and Control Position

What the leadership team is trying to achieve and protect

Outcome	Approved direction
Enterprise value	[Define growth, service, public impact, productivity, resilience, quality, safety, or decision outcomes.]
People	[Define how Alpha should amplify people, reduce friction, preserve judgment, and support adaptation.]
Trust	[Define customer, employee, citizen, partner, regulator, and public commitments.]
Control	We will retain appropriate control over mandate, knowledge, identity, authority, evidence, action, continuity, portability, and exit.
Sourcing	We may use proprietary, open, open-weight, managed, sovereign, and internally operated components when they meet the required control and evidence.
Autonomy	We will use the lowest autonomy level that delivers the value and raise it only through Evidence Gates.
Accountability	Named people remain accountable for mandates, residual risk, consequential decisions, exceptions, outcomes, and harm.

Control, Ownership and Operating Responsibility

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



Leadership chooses the required control position; ownership is one option.

04

Scope and Non-Goals

What Alpha may support and what it is not

In scope	Out of scope without separate approval
Enterprise search, explanation, analysis, drafting, planning, coordination, workflow support, approved automation, and bounded agent action.	Unrestricted access to the full enterprise corpus.
Approved personal and enterprise agents with explicit identity and delegation.	Implicit inheritance of a person's permissions by an agent.
Model and tool portfolios selected for work, data, risk, performance, and cost.	One universal model mandated for every task.
Governed use of Enterprise Knowledge and State.	Generated memory replacing authoritative records.
Evidence-based Delegated Autonomy for defined units of work.	Broad autonomous authority granted to Alpha or a vendor.
Controlled external partners and services.	Transfer of executive accountability or residual risk acceptance to a provider.

05

Governance and Accountability

Who decides, who owns, and who challenges

Role	Accountability
CEO / accountable executive	Approves mandate, enterprise control position, human-authority boundaries, material residual risk, and stop decisions.
Business process owner	Owns the work, customer or public promise, baseline, value, exceptions, and operating outcome.
Transformation leadership	Coordinates priorities, dependencies, capacity, Evidence Gates, and executive rhythm.
CIO / technology owner	Owns platform, integration, identity infrastructure, reliability, portability, operations, and technical risk within mandate.
Data owners and stewards	Own definitions, quality, access, provenance, current state, retention, and correction.
Risk, Security, Privacy, Legal, Compliance, Procurement	Provide specialist review and controls according to use, data, consequence, sector, and sourcing.
People / workforce leadership	Owns Workforce Charter, communication, skills, role impacts, employee rights, adoption, and appeal.
Independent assurance	Challenges evidence and controls without owning delivery success.

TIO governance rule

Governance is designed and owned in Leadership. Data, Technology, Factory, and People apply the rules in their work.

06

Human Authority

Decisions that remain human unless explicitly and legitimately delegated

- Enterprise strategy, mandate, risk posture, and allocation of accountability.
- Acceptance of material residual risk.
- Rights-affecting, safety-critical, legally reserved, highly ambiguous, or difficult-to-reverse decisions unless law and governance explicitly permit delegation.
- Changes to the autonomy level, action limits, data rights, or stop conditions of a material workflow.
- Exceptions outside the approved policy or evidence envelope.
- Employment, customer, citizen, supplier, financial, legal, or public-interest decisions where legitimate human judgment and due process are required.
- Approval to resume after a material incident or control failure.
- Final accountability for outcomes and harm.

Human approval must be meaningful. The approver receives clear information about the action, evidence, uncertainty, consequences, alternatives, and recovery. Deceptive, overloaded, or automatic approval is not human oversight.

07

Enterprise Knowledge and State

What Alpha may know and how it may know it

Mandate requirement	Approved position
Ownership	Priority domains have named business data and knowledge owners.
Classification	Information is classified by sensitivity, purpose, jurisdiction, retention, and AI-use rules.
Access	Retrieval is authorized by person, agent, purpose, source, fields, action, and duration.
Grounding	Alpha uses current, permitted, attributable sources; output identifies source and uncertainty where material.
Authoritative state	Financial, legal, identity, employment, inventory, safety, and other material state remains in systems of record.
Memory	Session, personal, team, workflow, and enterprise memory have explicit retention and correction rules.
Training and adaptation	Enterprise records are not used for training, fine-tuning, or provider improvement without explicit approval and rights.
Correction and deletion	Errors, outdated knowledge, access requests, retention expiry, and legal obligations have defined processes.

08

Models, Agents, Tools, and Sourcing

Approved portfolio and provider boundaries

Decision area	Mandate
Model portfolio	Use several approved models and deterministic methods rather than permanent monoculture.
Agent identity	Every agent has an owner, purpose, version, scopes, tools, data rights, expiry, monitoring, and disable path.
Tool registry	Only approved, owned, versioned, and security-reviewed tools may be used for enterprise action.
Provider data use	Contracts must state processing, retention, training, logs, support, incidents, location, and deletion.
Accounts and credentials	Organization-controlled accounts, keys, roles, and administrative access where practical.
Portability	Evaluations, prompts, policies, contracts, data connectors, logs, and transition rights remain controlled and documented.
Partner transfer	External work leaves the organization with knowledge, documentation, accounts, code, configurations, and decision capability.


09

Autonomy Posture

Progressive delegation with expiry

TIO Autonomy Ladder

Authority attaches to a defined unit of work, not to a model or agent in the abstract



Default posture	Mandate
Observe / Recommend	May be broadly approved for low-risk work with appropriate data and quality controls.
Prepare	Requires clear review, source visibility, and no silent external release.
Execute after approval	Requires action contract, explicit approval, transaction validation, evidence, and recovery.
Bounded autonomy	Requires completed Autonomy Envelope, Gate pass, limits, monitoring, circuit breaker, and accountable residual risk owner.
Domain autonomy	Requires sustained operating evidence, independent review, recovery testing, executive mandate, and periodic reauthorization.
Expiry and downgrade	Every Delegated Autonomy authorization has a review date and may be reduced or revoked when evidence, conditions, or risk change.

10

Action and Transaction

What Alpha may do

Action class	Default requirement
Read and retrieve	Purpose-bound access, minimization, provenance, audit.
Calculate and analyze	Approved methods, source data, validation, uncertainty, and reproducibility appropriate to consequence.
Draft and prepare	Human review before external, legal, financial, employment, or public commitment unless a separate Envelope permits otherwise.
Create or update records	Schema, policy, authority, transaction service, idempotency, evidence, and recovery.
Send, publish, pay, approve, deny, or control equipment	Privileged action with explicit limits, approval or bounded authority, monitoring, and stop path.
Agent-to-agent delegation	Preserve initiating identity, purpose, scopes, audience, time, and downstream authorization.
Emergency action	Only where explicitly designed, tested, and governed; immediate evidence and post-event review.

11

Assurance, Incidents, and Recovery

Evidence is part of the operating system

Required capability	Mandate
Evaluation	Use-case-specific acceptance suite and thresholds before launch and after material change.
Monitoring	Track quality, exceptions, policy denials, drift, cost, latency, incidents, and human overrides.
Audit	Retain enough identity, source, model, tool, approval, transaction, and outcome evidence to reconstruct material events.
Stop	Provide circuit breakers and authority revocation at workflow, agent, tool, model, and provider levels.
Recovery	Define rollback, compensation, reconciliation, state restoration, manual fallback, and continuity.
Incident response	Name business, technology, security, privacy, legal, communications, and vendor contacts.
Review	Material incidents, threshold breaches, and repeated overrides trigger an Evidence Gate and possible Delegated Autonomy downgrade.

12

People and Workforce Charter

Personal agents require trust, transparency, and support

Charter topic	Required decision
Purpose	How personal and enterprise agents help people and what they are not used for.
Visibility	What agents observe, store, share, infer, and report.
Monitoring	What is monitored for security and operations and what is prohibited as invisible surveillance.
Decision use	How AI information may influence performance, employment, scheduling, safety, or disciplinary decisions.
Correction and appeal	How people challenge data, outputs, actions, and automated decisions and obtain human support.
Skills and roles	What people must learn, which judgment remains human, and how work and careers change.
Accessibility	How tools support language, disability, literacy, device, and work context.
Continuity	How people work when Alpha is unavailable, wrong, or under incident control.

13

First 90 Days

Practical outputs

Time	Executive focus	Outputs
Days 1–30	Purpose, accountability, control, first workflow.	Approved mandate draft; accountable executive; initial human-authority boundaries; first workflow and baseline.
Days 31–60	Architecture and evidence.	Knowledge map; contracts; threat model; model and tool options; Autonomy Envelope; acceptance suite.
Days 61–90	Controlled operating evidence.	Limited release; monitoring; recovery rehearsal; Evidence Gate decision; next roadmap and control improvements.

Selection rule
The first workflow should matter enough to teach the architecture, yet be bounded enough that the organization can stop, recover, and learn.

14

Decision and Approval Record

Complete before mandate activation

Decision	Approved value / owner
Mandate owner	[Name and role]
Accountable executive	[Name and role]
Initial control position	[Shared / managed / dedicated / sovereign hosted / owned-operated by workflow]
Human-authority boundaries	[Summary]
Priority Enterprise Knowledge and State domains	[List]
Initial approved model, agent, and tool lanes	[List]
Initial prohibited actions	[List]
First workflow and autonomy level	[List]
Evidence Gate cadence	[Monthly / quarterly / event-triggered]
Residual risk acceptance authority	[Role]
Review and expiry date	[Date]
Approval	[CEO / accountable executive / required specialists]

15

Source Notes

Primary architecture and governance 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