



TIO

The Intelligent Organization

TIO End-State Architecture Blueprint

What the Intelligent Organization must control, how the architecture works, and how it can be built progressively

For CEOs, Executive Teams, CIOs, Transformation Leaders, Risk Leaders, and Enterprise Architects

Contents

01	Executive brief
02	Target state, status, and non-goals
03	Architecture at a glance
04	Governing principles
05	Business and operating model
06	Intelligence Core
07	Enterprise Knowledge and State
08	Identity and authority
09	Action and transaction
10	Trust and assurance
11	Personal and edge experience
12	Threat-to-control model
13	Non-functional requirements
14	Model and agent portfolio
15	Progressive realization path
16	TIO domain responsibilities
17	Architecture tests and first moves
18	Source notes

How to use this document

Read Sections 1 through 4 for the executive position. Use Sections 5 through 14 to understand the architecture. Use Sections 15 through 17 to plan, test, and govern practical movement toward the target state.

01

Executive Brief

One coherent enterprise intelligence experience, supported by disciplined architecture

TIO design hypothesis


Every enterprise should deliberately control the intelligence layer through which its people, knowledge, systems, decisions, and work interact. Ownership is one possible means of control, not the universal requirement.

Employees may increasingly work through personal or enterprise agents rather than navigating an application maze. Those agents can engage Alpha: the organization's complete Sovereign Enterprise Intelligence System. Alpha combines an Intelligence Core with Enterprise Knowledge and State, then connects both to identity, authority, controlled action, evidence, recovery, and accountable people.

The experience can become simpler than today. The architecture underneath must remain explicit. A probabilistic model cannot become the sole source of truth, inherit a person's authority automatically, or commit consequential transactions without deterministic controls.

TIO End-State Architecture

One coherent Alpha, two co-equal foundations, four operating planes



The TIO End-State Architecture at a glance.

Grounding principle

The model reasons. Repositories remember. Systems of record establish truth. Transaction services change state. People and governance establish legitimate authority.

Executive decision	Question
Purpose	What should enterprise intelligence help the organization become able to do?
Control position	What must the organization directly own, what may be sourced, and what control outcomes are non-negotiable?
Human authority	Which decisions and responsibilities remain with named people?
Knowledge and data	What may Alpha use, for what purpose, through which approved paths?
Autonomy	Which defined units of work may be observed, recommended, prepared, approved, or executed within bounds?
Action	Which authoritative services may commit changes, and what validation is required?
Assurance	What evidence, monitoring, challenge, stop, recovery, and review conditions apply?
People	How will roles, skills, trust, rights, appeals, and workforce communication change?

2026.08 ALIGNMENT | Three outcomes and transition connection

- The End-State Architecture is the destination for an organization that is more intelligent, more automated, and more in control through greater Enterprise Autonomy.
- The Work and Decision Map supplies the bounded business view. Transition Architecture defines how that bounded work moves from current to target operation.
- The Trust and Assurance plane uses the shared lifecycle: Register, Approve, Operate, Observe, Review, Change, Recover, Retire.

02

Target State, Status, and Non-Goals

*A reference architecture and controlled target-state hypothesis***Target-state statement**

The Intelligent Organization continuously senses, understands, decides, acts, learns, and adapts. It operates close to its current performance frontier while continually moving that frontier outward.

Version 1.2 is an enterprise reference architecture. It defines a credible destination, design principles, views, decisions, contracts, controls, and transition states. It becomes a target architecture for a specific organization only after adaptation to that organization's strategy, operating model, data, systems, people, risks, economics, and legal duties.

This architecture is	This architecture is not
A coherent enterprise target-state hypothesis	A universal product architecture
A control and accountability model	A claim that one model can run the enterprise
A reference for capability, data, technology, work, risk, and people decisions	A replacement for current-state architecture work
A framework for progressive Delegated Autonomy	A recommendation to maximize Delegated Autonomy everywhere
A sourcing-neutral control position	A demand to own all infrastructure
A basis for client-specific design and evidence	A claim of operational validation without implementation

The end state is not permanent peak optimization. An intelligent organization balances cost, speed, resilience, service, rights, trust, learning, and future adaptability. Short-term efficiency that destroys recoverability, capability, or legitimacy is not optimization.

03

Architecture at a Glance

Alpha has two co-equal foundations and four operating planes

Foundation	Role
Intelligence Core	Replaceable models, agents, reasoning methods, orchestration, evaluations, specialist methods, tools, and human expertise that understand, plan, coordinate, and learn.
Enterprise Knowledge and State	Governed knowledge, intellectual property, definitions, authoritative records, policies, evidence, relationships, events, and current operating condition.

Operating plane	Role
Identity and Authority	Identifies people, agents, models, tools, services, and devices and determines what each may see, infer, recommend, approve, or do.
Action and Transaction	Provides controlled tools, workflows, APIs, equipment interfaces, and authoritative services through which work changes state.
Trust and Assurance	Evaluates, constrains, monitors, records, challenges, stops, recovers, and improves Alpha and the work it coordinates.
Personal and Edge Experience	Gives people context-aware interfaces while preserving enterprise authority, privacy, resilience, and visibility.

Architectural boundary

The two foundations create the enterprise capability. The four operating planes make Alpha safe, executable, governable, and resilient.

04

Governing Principles

Design rules that survive model and vendor changes

Principle	Meaning
Enterprise first	Begin with strategy, work, value, trust, and people. Technology choices follow.
Control over slogans	Prove mandate, knowledge, authority, evidence, action, recovery, and exit control.
Ground, do not indiscriminately train	Retrieve permitted, current, attributable information from governed sources.
Authoritative state remains authoritative	Generated memory or model output never silently replaces systems of record.
Every actor has identity	A person, agent, model, tool, service, and device are not the same principal.
Authority is purpose-bound	Delegation is narrow, short-lived, reviewable, revocable, and recorded.
Commit deterministically	Consequential state changes are validated and committed by authoritative services.
Delegated Autonomy is earned by evidence	Authority expands only when performance, controls, adoption, and recovery are proven.
Design for replacement	Models, providers, and infrastructure should change without destroying the business capability.
Resilience is sovereignty	The organization must be able to continue, degrade safely, stop, recover, and change direction.
People remain accountable	Named people own mandates, residual risk, consequential decisions, exceptions, and harm.
Simplicity at the surface, discipline underneath	A simple experience must not erase control boundaries.

05

Business and Operating Model

Alpha changes how work is requested, coordinated, decided, and improved

The business view begins with outcomes and units of work. It does not begin with the model portfolio. Each priority workflow has a process owner, accountable executive, users, customer or public promise, baseline, value hypothesis, authority model, evidence requirements, and transition path.

Business element	Target-state design
Enterprise direction	Leadership defines outcomes, risk posture, human authority, sourcing position, and what the organization will not delegate.
Value streams	Alpha supports or coordinates bounded work across customer, employee, operational, regulatory, and management value streams.
Decision rights	People decide mandates, residual risk, consequential exceptions, and changes to Delegated Autonomy.
Human and machine roles	Machines sense, analyze, prepare, coordinate, and execute within bounds; people judge, authorize, challenge, care, negotiate, and remain accountable.
Operating rhythm	Portfolio, Evidence Gate, incident, value, and workforce reviews connect architecture to management decisions.
Learning loop	Outcomes update knowledge, evaluations, instructions, controls, training, and work design through governed review.

Operating rule

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

06

Intelligence Core

A replaceable portfolio, not permanent model monoculture

Capability	Plain-English test
Understand intent	Can the Intelligence Core interpret the mandate, role, work, requested outcome, and relevant constraints?
Reason and plan	Can it choose methods, identify assumptions and uncertainty, and propose a controlled plan?
Route intelligence	Can it select the right model, solver, agent, tool, or human specialist for the task?
Use enterprise knowledge	Can it retrieve permitted, current, attributable context without treating generated text as truth?
Coordinate work	Can it manage steps, dependencies, approvals, tools, agents, and follow-up within a defined workflow?
Verify	Can it challenge outputs using rules, independent models, test cases, calculations, or human review?
Learn safely	Can the organization improve instructions, routing, evaluations, and work design without uncontrolled self-modification?
Remain replaceable	Can a model or provider change without rebuilding the enterprise capability?

The Intelligence Core may use proprietary frontier models, open or open-weight models, smaller task-specific models, deterministic solvers, search, rules, analytics, vision, speech, robotics, and human specialists. A single logical Alpha can present one identity while routing work across many physical forms of intelligence.

Routing factor	Why it matters
Required capability	Use frontier capability only where the work needs it.
Information sensitivity	Keep restricted work in approved environments and minimize exposed context.
Reliability and verification	Choose methods that can meet the workflow's consistency and evidence needs.
Latency and availability	Match model and hosting choices to operating deadlines and continuity requirements.
Cost and utilization	Avoid paying frontier-model prices for routine classification, extraction, or rules.
Portability	Prefer contracts and abstractions that keep prompts, tools, evaluations, and business logic movable.

07


Enterprise Knowledge and State

The crown-jewel foundation is built progressively

Enterprise Knowledge and State includes more than data. It holds what the organization knows, owns, has learned, and knows to be true now. It may include documents, master data, transaction records, policies, definitions, knowledge graphs, event streams, cases, sensor readings, current approvals, lineage, provenance, and evidence.

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.

A partial-credit path from one priority workflow to shared enterprise current state.

Failure mode	What happens
No business owner	Technical teams cannot resolve meaning, access, quality, or acceptable use.
Stale or conflicting truth	Alpha produces plausible answers from obsolete or contradictory sources.
Broad undifferentiated access	Sensitive information is exposed beyond purpose or authority.
Missing provenance	Leaders cannot trace a recommendation, action, or number to its source.
Fragile data flow	Automations and agents fail when upstream fields, timing, formats, or systems change.
Generated memory becomes authoritative	Unverified output displaces legal, financial, employment, identity, inventory, or safety records.
No retention and deletion model	The organization cannot meet records, privacy, legal hold, or correction obligations.
No current-state recovery	The organization cannot restore or reconcile state after failed or duplicated actions.

Required contract field	Example
Identity and purpose	Who is asking, which agent acts, and why the information is needed.
Permitted sources and fields	Approved repositories, records, classifications, and field-level minimization.
Freshness and time	As-of time, event time, expiry, and acceptable staleness.
Provenance	Source, owner, version, retrieval time, transformation, and citations.
Permitted inference	What may be inferred or combined; prohibited sensitive or rights-affecting inference.
Retention and memory	Whether the context may be cached, remembered, trained on, or deleted.
Response schema	Required fields, confidence, uncertainty, citations, and error conditions.
Write-back	Whether and how approved corrections, evidence, or state may be written.

08

Identity and Authority

No implicit inheritance of human rights by agents

A user may be allowed to view a record while the user's personal agent is not allowed to store it, infer from it, or share it with another agent. Identity and authority must be evaluated per request, purpose, information, action, device, time, and risk condition.

Actor	Identity and authority requirement
Person	Verified role, employment or customer relationship, authentication, and current access.
Personal agent	Separate identity, device binding, explicit delegation, narrow scopes, visible activity, revocation.
Enterprise agent	Named owner, approved purpose, workflow scope, tool and data rights, expiry, monitoring.
Model	Approved version, provider, environment, data-use conditions, evaluation status, and routing policy.
Tool or service	Registered capability, schema, owner, version, authorization, integrity, and allowed operations.
Device	Security posture, location or network context where relevant, and approved connection path.
Human approver	Legitimate decision right, segregation of duties, evidence visibility, and non-repudiation.

- No implicit trust based on network location, seniority, device ownership, or prior approval.
- Least privilege and purpose-bound access.
- Short-lived, revocable delegation.
- Audience-bound tokens and no uncontrolled token passthrough.
- Clear delegation chains for agent-to-agent work.
- Separation of preparation, approval, and commit for consequential actions.
- Audit records that establish who or what acted, for whom, under which authority, using which evidence.


09

Action and Transaction

Interpret, plan, validate, approve, commit, and record

Controlled Action Path

Intelligence interprets and prepares; authoritative services validate and commit

*The deterministic commit boundary keeps probabilistic reasoning separate from authoritative state.*

Action contract field	Required decision
Operation	The exact action or command allowed.
Input schema	Required and optional fields, formats, validation, and source.
Authority	Which identity, purpose, role, amount, case, or context permits the call.
Limits	Financial, operational, rate, time, geographic, and volume bounds.
Approval	When explicit human, specialist, or second-agent approval is mandatory.
Idempotency	How retries avoid duplicate or conflicting actions.
Transaction boundary	What commits atomically and what requires compensation or reconciliation.
Output and evidence	Authoritative result, record ID, timestamp, approver, model/tool versions, and outcome.
Failure and recovery	Timeout, retry, rollback, compensation, manual recovery, and escalation.

Financial ledgers, payment release, payroll, identity, access control, legal records, inventory, benefits, employment status, safety interlocks, and other material states require deterministic validation and durable records. AI may prepare and coordinate; the authoritative service commits.

10

Trust and Assurance

Continuous evidence, challenge, incidents, recovery, and improvement

Assurance capability	Evidence
Evaluation	Task-specific test sets, acceptance thresholds, edge cases, subgroup analysis where relevant, and regression tests.
Runtime monitoring	Quality, latency, cost, tool use, access, exceptions, drift, incidents, and human overrides.
Policy enforcement	Identity, data, tool, action, amount, purpose, and Delegated Autonomy rules applied before execution.
Independent challenge	Rules, secondary models, calculations, specialists, or human review for material outputs.
Audit and records	Prompt or instruction context where appropriate, sources, model and tool versions, delegations, approvals, transactions, and outcomes.
Incident response	Containment, authority revocation, provider escalation, customer or employee correction, and regulatory response where required.
Recovery	Rollback, compensation, state restoration, failover, degraded operation, and continuity.
Learning	Approved changes to work design, data, instructions, routing, evaluations, training, and governance.

Sovereignty test

An organization that cannot stop, recover, investigate, and change providers does not control its intelligence layer.

11

Personal and Edge Experience

A simple interface with clear boundaries and employee protections

Edge capability	Boundary
Personal context	May understand role, commitments, preferences, and local working context within approved privacy boundaries.
Local processing	May perform appropriate summarization, drafting, classification, or authentication on the device.
Enterprise request	Must authenticate the person, agent, device, purpose, and requested authority.
Agent representation	May coordinate with enterprise agents only within explicit delegation and visible scope.
Offline work	May support local tasks, but authoritative state changes wait for controlled enterprise services.
Monitoring	The person should know what the agent observes, stores, shares, infers, and reports.
Correction and appeal	Employees and affected people need understandable correction, appeal, and human-support paths.
Continuity	People must retain the ability to work, recover, and make decisions when Alpha is unavailable or wrong.

12

Threat-to-Control Model

Named threats, architecture answers, and recovery

Threat-to-Control Model

Agentic systems need controls that answer named failure modes

Threat	Failure mode	Prevent	Detect	Respond / recover
Prompt or context injection	Untrusted content changes agent intent	Separate instructions from data; sanitize context; allowlisted tools; policy checks	Detect anomalous prompts and tool plans; log source and provenance	Pause workflow; revoke delegation; remove poisoned context
Tool poisoning or misuse	A tool, schema or description causes unsafe action	Signed registry; version pinning; narrow schemas; least privilege; sandboxing	Tool integrity checks; runtime policy; unusual-call monitoring	Disable tool, rotate credentials, reconcile affected transactions
Confused deputy	An agent acts with authority borrowed from another actor	Distinct identities; audience-bound tokens; per-client consent; purpose-bound delegation	Authorization telemetry; delegation chain review	Revoke chain, investigate, reauthorize explicitly
Retrieval or memory poisoning	False or malicious content enters working memory	Source allowlists; provenance; write controls; quarantine; human review	Freshness, anomaly and source-reputation checks	Remove entry, re-index, review dependent decisions
Excessive agency	Agent can execute more than the work requires	Autonomy Envelope; action limits; segregation of duties; approval thresholds	Exception rates, blast-radius monitoring, circuit breakers	Stop, rollback or compensate, downgrade autonomy
Cascading or partial failure	Multi-step action leaves inconsistent state	Idempotency; transaction boundaries; staged commit; compensation design	End-to-end tracing; reconciliation; timeout and retry limits	Compensate, restore authoritative state, reopen for human review
Supply-chain compromise	Model, code, plugin or provider is compromised	Due diligence; signed artifacts; isolation; SBOM; alternative provider path	Integrity monitoring; version drift; threat intelligence	Isolate, fail over, restore known-good version
Sensitive-information disclosure	Agent exposes data beyond purpose or authority	Data minimization; field-level access; output filtering; no implicit memory	DLP, access logs, canary data, exfiltration alerts	Contain, notify, delete, investigate and correct permissions

Version 1.2 connects agentic threats to preventive, detective, response, and recovery controls.

The confused-deputy problem deserves special attention. An agent or proxy must not use its own standing authority to satisfy a request that the initiating client, agent, or person was not entitled to make. Delegation chains, token audience, per-client consent, purpose, and downstream authorization must remain explicit.

Threat	Primary architecture constraints
Prompt or context injection	Treat retrieved content as data, separate instructions, sanitize context, constrain tools, apply policy before action.
Tool poisoning or misuse	Signed registry, owner and version, schema validation, least privilege, sandboxing, runtime policy, disable path.
Confused deputy	Distinct identities, audience-bound tokens, per-client consent, purpose-bound delegation, downstream authorization.
Retrieval or memory poisoning	Source allowlists, provenance, write controls, quarantine, freshness, anomaly checks, dependency review.
Excessive agency	Autonomy Envelope, limits, approvals, segregation, monitoring, circuit breakers, downgrade and recovery.

Threat	Primary architecture constraints
Cascading or partial failure	Idempotency, staged commit, transaction boundaries, compensation, tracing, reconciliation.
Supply-chain compromise	Due diligence, signed artifacts, isolation, version control, integrity monitoring, failover and known-good recovery.
Sensitive-information disclosure	Minimization, field-level access, retention rules, output controls, DLP, audit, containment and correction.

13

Non-Functional Requirements

The architecture must work under real operating constraints

Requirement	Design question	Illustrative evidence
Security	Are identities, delegation, tools, data, actions, secrets, and supply chains protected?	Threat model, control test, access review, incident exercise.
Privacy and records	Can purpose, minimization, retention, deletion, access, correction, residency, and legal holds be enforced?	Data contract, privacy review, records schedule, deletion test.
Reliability	Does the workflow meet quality and consistency thresholds under normal and perturbed conditions?	Representative test set, exception rate, regression results.
Latency	Can the system respond within the work's decision window?	P50/P95/P99 latency by workflow step.
Availability	Can mission-essential work continue when models, data, networks, or providers fail?	SLOs, failover test, degraded operating procedure.
Recoverability	Can failed actions be stopped, rolled back, compensated, reconciled, and investigated?	Recovery time and point objectives; completed simulation.
Auditability	Can a material outcome be reconstructed from identity, source, model, tool, approval, transaction, and evidence records?	End-to-end trace and evidence package.
Portability	Can the organization move models, providers, accounts, configurations, evaluations, and data without losing the capability?	Migration rehearsal and acceptance-test comparison.
Operability	Can named teams monitor, support, update, patch, and improve the service?	Runbooks, staffing, on-call model, training, support metrics.
Cost	Are fixed and variable costs visible and connected to value and service levels?	Whole-life cost model, unit economics, utilization, budget limits.
Accessibility and usability	Can intended people use, understand, challenge, and recover from the system?	User testing, accessibility review, adoption and override evidence.

14

Model and Agent Portfolio

One Alpha can use many interchangeable forms of intelligence

Model replaceability is not created by declaring models replaceable. It requires stable contracts, organization-owned evaluation sets, separation of business logic and data connectors from model-specific behaviour, documented dependencies, and periodic alternative testing.

Portability commitment	Test
Organization-owned evaluations	Run the same acceptance suite against at least one credible alternative.
Stable tool and response schemas	Confirm a replacement can use the required tools and produce valid responses.
Separated instructions and business logic	Keep policies, workflow rules, prompts, retrieval, and transactions outside proprietary model internals where practical.
Export and transition rights	Verify data, configuration, logs, fine-tunes, adapters, and account rights before contracting.
Known switching costs	Record performance, integration, revalidation, training, and operational costs.
Provider failure path	Test how the workflow degrades, fails over, pauses, or returns to human operation.

Open and open-weight models can strengthen bargaining power, deployment choice, customization, and local control. They also transfer operating responsibility for security, updates, serving, evaluation, talent, and continuity. The sourcing decision should balance capability, trust, cost, control, and learning.


15

Progressive Realization Path

Build bounded operating slices, not an abstract platform

Progressive Realization Path

The end state is built through bounded operating slices, not a single transformation project



Progress is uneven by design. An organization may operate at advanced autonomy in one bounded domain while remaining at “recommend” or “prepare” for rights-affecting decisions.

Horizon	Typical outcomes	Evidence to move forward
1. Governed access and visibility	Inventory uses, owners, data lanes, approved tools, decision rights, and baseline work.	Named owners, usage visibility, safe lanes, risk triage.
2. Priority grounding and controlled tools	One workflow uses attributable knowledge and narrow tool access.	Data and tool contracts, quality baseline, user feedback.
3. Model and agent orchestration	Routing, evaluations, specialist agents, and controlled registry operate.	Acceptance suite, access logs, cost and performance evidence.
4. Bounded autonomy	Selected actions run inside an approved Envelope.	Gate pass, thresholds met, monitoring and recovery tested.
5. Cross-domain intelligence and current state	Priority domains share identity, semantics, events, and assurance.	Lineage, interoperability, reconciliation and continuity evidence.
6. Domain autonomy	A defined domain operates within an executive mandate and continuous assurance.	Sustained operating evidence, incidents, recovery, residual risk review.

Progress will be uneven. A company may reach bounded autonomy in infrastructure monitoring while remaining at “recommend” for employment, credit, health, public-benefit, or other rights-affecting decisions.

16

TIO Domain Responsibilities

The five domains build and operate the architecture together

Domain	Primary responsibilities
Leadership	Enterprise direction, Alpha mandate, control position, decision rights, risk posture, human authority, residual risk, governance, and transformation rhythm.
Data	Ownership, quality, access, flow, insight, provenance, definitions, authoritative records, current state, retention, and correction.
Technology	Model and agent portfolio, identity infrastructure, integration, hosting, observability, resilience, portability, security, and operations.
Factory	Discovery, work redesign, delivery, evaluation, Evidence Gates, scaling, stopping, and value measurement for bounded workflow slices.
People	Workforce Charter, roles, skills, communication, adoption, accessibility, employee rights, challenge, appeal, and human capability.

Leadership designs governance. Data, Technology, Factory, and People deploy and manage the rules within their work. Alpha does not become a sixth TIO domain; it is an enterprise architecture and operating capability created across the five domains.

17

Architecture Tests and First Moves

How leaders can begin without overbuilding

Architecture test	Question
Control test	Can the organization explain and prove what it controls across mandate, knowledge, authority, action, evidence, recovery, and exit?
Truth test	Can every material fact or state be traced to an authoritative source?
Authority test	Can every material action be traced through a legitimate delegation chain?
Commit test	Does an authoritative service validate and record consequential state changes?
Autonomy test	Is machine authority bounded by unit of work, thresholds, monitoring, stop, recovery, and owner?
Threat test	Does each named threat have preventive, detective, response, and recovery controls?
Portability test	Can the capability survive a model or provider change?
People test	Can affected people understand, challenge, correct, appeal, and continue working?
Evidence test	Can leaders see value, quality, risk, incidents, adoption, and recovery evidence?

1. Approve an Enterprise Intelligence Mandate and control position.
2. Choose one bounded workflow that matters, has an owner, and can be measured.
3. Map the knowledge, authoritative state, identities, tools, approvals, and transaction path for that workflow.
4. Set the initial autonomy level and complete an Autonomy Envelope before building.
5. Define evaluation thresholds, monitoring, stop conditions, and recovery before launch.
6. Run the workflow through Evidence Gates and decide whether to expand, redesign, continue, reduce, or stop.
7. Use the evidence to improve the architecture and build reusable patterns for the next workflow.

18

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