Delegated Autonomy and Evidence Framework

How the organization delegates machine authority, proves that it works, and takes it back when conditions change

For accountable executives, business owners, Risk, Technology, Factory, Legal, Privacy, Security, Audit, and Enterprise Architecture

Document Version 1.3 FINAL | TIO-KR-2026.08.0 | August 2026

Contents

How to use this document

Use Sections 1 through 4 to agree on the operating model. Use Sections 5 and 6 as the mandatory design and decision instruments. Use Sections 7 through 9 as worked examples. Use Sections 10 through 13 to operate and mature the capability.

01 Executive Brief

Delegated Autonomy is delegated authority for a defined unit of work

TIO position
The organization should seek the right delegated autonomy for the work, supported by the right evidence, controls, accountability, and recovery. It should not seek the greatest delegated autonomy technically possible.

Enterprise AI becomes consequential when it can change records, send communications, release money, operate equipment, affect people, or trigger other systems. At that point, the issue is no longer whether the model can produce a plausible answer. The issue is whether the organization has deliberately delegated authority and can prove that the work remains safe, useful, lawful, reversible, and accountable.

Inside Alpha, the Intelligence Core may interpret, recommend, plan, and coordinate. Enterprise Knowledge and State provides governed context, authoritative records, and current operating condition. Delegated Autonomy decisions specify how those foundations may be used and which authoritative services may commit a change.

The TIO Delegated Autonomy Ladder. Authority increases only when evidence and control increase with it.

Authority and Scope

Executive question	Required answer
What may act?	The named person, personal agent, enterprise agent, model, tool, service, and device identities.
For whom?	The principal or accountable owner and the valid delegation chain.
On what work?	A bounded task, workflow step, decision, record, transaction, or operating domain.
Within which limits?	Data, action, amount, population, time, location, consequence, confidence, and approval limits.

Evidence, Change, and Accountability

Executive question	Required answer
What proves fitness?	Baseline, acceptance thresholds, independent checks, runtime evidence, and comparison to alternatives.
What happens when conditions change?	Escalation, pause, downgrade, rollback, compensation, recovery, and human takeover.
Who remains accountable?	A named executive or process owner who can accept residual risk and stop the work.

Operating rule
Delegated Autonomy is valid only while the evidence, operating conditions, and accountable ownership remain within the approved Envelope.

2026.08 ALIGNMENT | Terminology boundary

- Enterprise Autonomy is an organizational outcome that TIO protects.
- Delegated Autonomy is machine or system authority granted to a defined unit of work.
- This framework governs Delegated Autonomy. The Ladder, Envelope, Evidence Gates, stop conditions, and recovery rules remain unchanged.

02 What Delegated Autonomy Means in TIO

Authority attaches to work, not to a model or agent

Delegated Autonomy is the organization's explicit delegation of authority to perform a defined unit of work without a person approving every instance. The same agent may operate at different levels inside one workflow. It may observe one step, recommend another, prepare a transaction, execute a low-value update within limits, and stop for human approval before a consequential decision.

Boundary rule
A system is never "Level 5." A unit of work may be assigned Level 0 through Level 5 under a current Delegated Autonomy Envelope.

Delegated Autonomy is assigned to	Examples
A task	Classify an incoming document.
A workflow step	Prepare a refund calculation.
A decision	Recommend a maintenance priority.
An action	Update a delivery date within an approved range.
A transaction	Release a payment below a defined amount after deterministic validation.
An operating domain	Balance energy demand within a tested control envelope.

Named people remain accountable for the mandate, design, delegation, residual risk, exceptions, outcomes, and harm. An AI system may assist judgment or execute bounded work. It does not absorb executive accountability.

03 The Delegated Autonomy Ladder

Six levels of progressively delegated machine authority

Level	Machine role	Human role
0 - Observe	Read, classify, monitor, detect, and summarize.	A person interprets and acts.
1 - Recommend	Generate options, predictions, or recommended decisions.	A person decides and acts.
2 - Prepare	Draft the work, transaction, communication, or plan.	A person reviews, changes, and releases.
3 - Execute after approval	Execute the specific prepared action after explicit approval.	A person approves each consequential instance.
4 - Bounded delegated autonomy	Execute within a defined envelope, thresholds, time period, and escalation conditions.	People monitor exceptions, review evidence, and remain accountable.
5 - Domain delegated autonomy	Manage a defined operating domain through continuous sensing, action, monitoring, and recovery.	People set outcomes, constraints, evidence, intervention rights, and stop conditions.

The ladder is not a maturity race. Level 1 can be the correct permanent design for a rights-affecting decision. Level 4 can be appropriate for a high-volume, reversible, tightly constrained operating task. Level 5 should be rare and limited to domains with strong observability, deterministic controls, tested recovery, and accountable oversight.

04 Entry and Exit Criteria by Level

Evidence and control must rise with delegated authority

Level	Minimum entry evidence	Mandatory downgrade / stop triggers
0 - Observe	Approved read-only sources; defined purpose; access and logging; no write path.	Material data exposure, unreliable sensing, or policy breach.
1 - Recommend	Representative quality tests; uncertainty visible; source evidence; human decision retained; bias and impact review where relevant.	Recommendation quality falls below threshold, users over-rely, or affected groups experience material harm.
2 - Prepare	All Level 1 evidence plus deterministic validation of drafts, clear review interface, no hidden release path, and reviewer competence.	Review becomes superficial, preparation errors exceed tolerance, or drafts conceal material uncertainty.
3 - Execute after approval	Identity and delegation; explicit approval; action preview; business-rule checks; transaction ID; rollback or compensation; complete record.	Approval is bypassed, deceptive, stale, or not attributable; action cannot be reconstructed or recovered.
4 - Bounded delegated autonomy	All Level 3 controls plus tested thresholds, continuous monitoring, exception routing, rate and value limits, segregation of duties, circuit breakers, and periodic Evidence Gates.	Threshold breach, drift, incident, control failure, data staleness, unexplained exception growth, or recovery failure.
5 - Domain delegated autonomy	Proven Level 4 operation over a defined period; redundant sensing and control; independent assurance; resilience tests; manual or alternate operation; executive risk acceptance.	Loss of observability, compounding error, material environmental change, unresolved incident, or inability to restore safe operation.

Promotion rule
A higher level requires more than improved model accuracy. It requires better identity, data, action control, monitoring, recoverability, operating capability, and accountability.

05 The Delegated Autonomy Envelope

The controlled delegation record for one unit of work

Envelope field	Required content
1. Unit of work	The exact task, decision, action, transaction, or domain covered.
2. Intended outcome	Business, customer, public, workforce, risk, or operational result.
3. Accountable owner	Named executive or process owner who owns the delegation and residual risk.
4. Level and duration	Approved delegated autonomy level, effective date, expiry, and review cadence.
5. Identities and delegation	People, agents, models, tools, services, devices, and authority chain.
6. Knowledge and data	Approved sources, classifications, purpose, prohibited data, freshness, provenance, retention, and inference limits.
7. Permitted actions	Allowlisted tools and actions, value limits, populations, locations, channels, and transaction bounds.
8. Prohibited actions	Decisions, inferences, records, communications, or transactions outside the mandate.
9. Evidence and thresholds	Baseline, acceptance criteria, confidence, quality, fairness, safety, cost, latency, and recovery requirements.
10. Monitoring and escalation	Runtime measures, exception routes, human takeover, alert severity, and review frequency.
11. Stop and recovery	Circuit breakers, rollback, compensation, reconciliation, continuity, and incident response.
12. Approval and residual risk	Specialist inputs, Evidence Gate decision, accountable approval, conditions, and expiry.

The Envelope is a living control record. A material model, data, tool, policy, interface, user-population, or operating-context change reopens the Envelope and requires targeted revalidation.

06 Evidence Gates

Six decisions from idea to continued authority

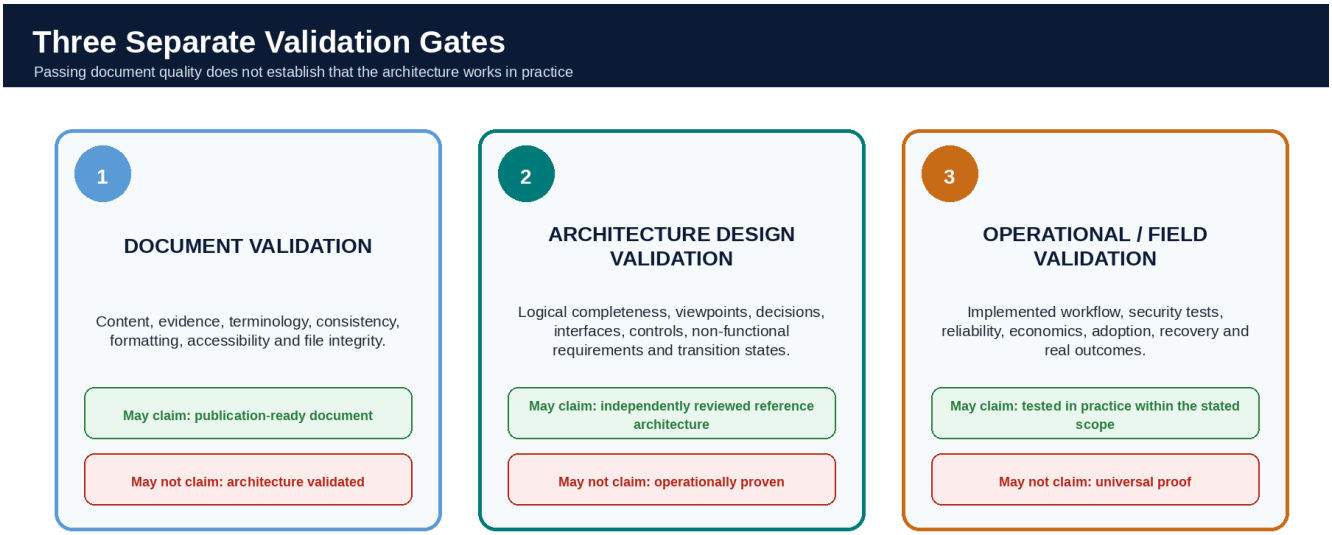

Evidence Gates turn delegated autonomy into a sequence of accountable decisions.

Gate	Decision question	Possible decision
Gate 1 - Purpose and ownership	Is the problem worth solving? Is the work and accountable owner clear?	Explore, redefine, or stop.
Gate 2 - Work and risk design	Has the workflow been redesigned? Are affected people, rights, risks, and controls visible?	Design, narrow, or stop.
Gate 3 - Build readiness	Are knowledge, data, technology, identity, integration, talent, and support ready enough?	Build, remediate, or wait.
Gate 4 - Release readiness	Has the proposed level passed representative, adversarial, access, action, recovery, and user tests?	Release within limits, redesign, or stop.
Gate 5 - Operating evidence	Does runtime performance meet the approved thresholds compared with the baseline?	Continue, expand, constrain, downgrade, or stop.
Gate 6 - Renewal or retirement	Does the delegation still fit strategy, law, risk, technology, users, and operating context?	Renew, change, transition, retire, or stop.

Evidence Gates are management decisions, not technical test reports. The decision record should show the evidence considered, dissent, assumptions, residual risk, conditions, expiry, and named approval.

07 Worked Delegated Autonomy Envelope

Illustrative example: invoice coding and low-value exception routing

Use of numbers
Illustrative thresholds only. Each organization must set thresholds based on the work, consequence, legal duties, baseline performance, risk posture, and recovery capability.

Envelope field	Completed example
Unit of work	Read supplier invoices, extract fields, assign accounting code, match to purchase order, and route exceptions. The system may post a matched invoice only when all Level 4 conditions are met.
Intended outcome	Reduce manual handling time while preserving payment accuracy, segregation of duties, audit evidence, and supplier experience.
Accountable owner	Vice President, Finance Operations. Residual risk accepted jointly with Controller.
Level and duration	Level 4 bounded delegated autonomy for 90 days. Automatic expiry on 31 October 2026 unless Gate 5 renews it.
Identities and delegation	Enterprise invoice agent with unique workload identity; approved document-extraction model; finance coding service; ERP posting service. No shared human credentials or token passthrough.
Knowledge and data	Approved supplier master, chart of accounts, purchase orders, goods receipts, tax rules, invoice images, and prior validated coding examples. No employee banking data or unrelated contracts. Source provenance and as-of time returned with every decision.
Permitted actions	Post invoices at or below CAD $2,500 only when supplier, purchase order, receipt, tax, currency, bank account, amount, and duplicate checks all pass. Maximum 250 postings per hour and CAD $100,000 per day.
Prohibited actions	No new supplier creation, bank-account change, non-purchase-order invoice posting, split invoice, manual journal entry, payment release, or action on sanctions or fraud alerts.
Evidence thresholds	Field extraction >=99.5%; coding accuracy >=98.0%; false auto-post <=0.1%; duplicate prevention 100% in test; 100% deterministic rule pass; median processing <=20 seconds; cost <=CAD $0.35 per invoice; no material control or privacy incident.
Monitoring and escalation	Daily exception and auto-post reconciliation; weekly sample of 200 auto-posts reviewed independently; alert at extraction <99.0%, coding <97.5%, exception rate >12%, stale master data >15 minutes, or any unauthorized tool request.
Stop and recovery	Immediate circuit breaker for bank mismatch, duplicate-control failure, unauthorized tool call, unexplained posting, or audit-log loss. Reverse or hold affected postings, reconcile all transactions since last known-good checkpoint, rotate credentials, and operate manually.
Approval and residual risk	Gate 4 approved by Controller, Finance Operations, Security, Privacy, Technology Service Owner, and Internal Controls. Residual risk limited to reversible misclassification below CAD $2,500 under continuous reconciliation.

08 Worked Evidence Gate Decision

Illustrative Gate 5: continue with tighter limits

Decision-record field	Evidence and decision
Operating period	1 August to 30 September 2026; 18,420 invoices received; 11,860 met auto-post eligibility.
Baseline	Median manual handling 7.8 minutes; coding accuracy 97.1%; average processing cost CAD $2.42; 3.2% invoices required rework.
Observed performance	Extraction 99.72%; coding 98.43%; false auto-post 0.08% (9 invoices); duplicate prevention 100%; median processing 13 seconds; average processing cost CAD $0.31; rework 1.1%.
Control evidence	All auto-posts passed deterministic rules. Two stale-data alerts stopped posting automatically. No unauthorized tool call, payment release, privacy incident, or audit-log loss.
Recovery evidence	Nine coding errors reversed within one business day; no incorrect payment. Circuit-breaker rehearsal restored manual operation in 11 minutes against a 15-minute requirement.
Human and supplier impact	Finance team reported lower repetitive work and increased exception complexity. Supplier inquiries fell 18%. Reviewers identified fatigue during the Friday sample review.
Dissent / concern	Internal Audit requested a lower daily value limit until the independent sample has covered quarter-end invoices.
Decision	Continue Level 4 for 90 days. Keep CAD $2,500 per invoice; reduce daily cap from CAD $100,000 to CAD $75,000 through quarter-end; add a second reviewer for Friday sampling; do not expand to non-purchase-order invoices.
Residual risk owner	Vice President, Finance Operations. Next Gate 5 review: 15 December 2026.
Automatic downgrade triggers	Any false auto-post above 0.15% over a rolling 7 days, duplicate-control failure, material incident, recovery >15 minutes, or two consecutive weekly coding samples below 98.0%.

Evidence-based movement
The decision did not simply “scale.” It preserved the delegated autonomy level, tightened one limit, improved a human control, and refused a higher-risk scope expansion.

09 End-to-End Workflow Example

One workflow can contain six delegated autonomy levels

Workflow step	Level	Operating design
Monitor incoming invoice channel	0 - Observe	Detect new files and summarize volume; no write access.
Identify likely duplicate or fraud pattern	1 - Recommend	Flag and explain; finance or fraud owner decides.
Extract invoice fields and prepare coding	2 - Prepare	Draft transaction and show source evidence.
Create exception case after reviewer approval	3 - Execute after approval	Reviewer approves the prepared case creation.
Post fully matched invoice below limits	4 - Bounded delegated autonomy	Execute only inside the approved Envelope and deterministic checks.
Manage the entire accounts-payable domain	5 - Domain delegated autonomy	Not approved. Payment release, supplier changes, fraud decisions, and material exceptions remain human-controlled.

This mixed design is usually stronger than choosing one delegated autonomy level for an application. It delegates routine work where evidence is strong and preserves human authority where consequences, ambiguity, rights, or recovery demand it.

10 Monitoring, Downgrade, and Recovery

Authority must be continuously renewable and reversible

Control	Minimum practice
Runtime measures	Quality, confidence, exceptions, drift, freshness, policy denials, tool calls, cost, latency, user overrides, affected-person impact, and incidents.
Independent review	Sampling or control testing sufficiently separate from the delivery team and the autonomous action.
Change detection	Model, prompt, policy, data, tool, provider, integration, interface, population, and context changes trigger revalidation.
Downgrade path	Level 5 to 4, 4 to 3, 3 to 2, or complete stop can be enacted without redesigning the whole service.
Recovery	Rollback, compensation, reconciliation, restore, manual fallback, communication, records, and lessons.
Expiry	Delegation expires automatically unless a named owner renews it based on current evidence.

Control	Minimum practice
Challenge and appeal	Employees, customers, residents, suppliers, and specialists can challenge material outcomes and reach a person.

A stop condition is not evidence that the design failed. A well-designed autonomous system detects when its authority is no longer justified and moves safely to a lower level or human operation.

11 Portfolio and Governance

See delegated autonomy across the enterprise, not one use case at a time

Portfolio field	What leaders need to see
Unit of work	Specific workflow step, action, or domain.
Business owner	Named accountable executive and process owner.
Current and target level	Level today, desired level, and reason.
Consequence tier	Impact on money, safety, rights, customers, employees, operations, and trust.
Evidence status	Last Gate, thresholds, exceptions, incidents, and expiry.
Dependencies	Knowledge, data, identity, tools, systems, people, vendors, and recovery.
Stop authority	Who can pause, revoke, downgrade, or retire the delegated autonomy.
Value evidence	What changed compared with the baseline and what value was realized.

Leadership owns the governance design, decision rights, risk posture, and accountability. Data, Technology, Factory, and People apply those rules in their own work. The Factory manages the use-case and improvement path; it does not accept residual risk on behalf of accountable executives.

12 Maturity Guidance

From implicit automation to managed delegated autonomy

Level	Observable maturity
0 - Absent	The organization cannot identify where AI or automation acts or who is accountable.
1 - Initial / Ad-Hoc	Some automated or agentic work exists, but authority, evidence, monitoring, and stop conditions depend on individuals.
2 - Repeatable	Basic Delegated Autonomy Envelopes and approvals are used for selected work more than once.
3 - Defined	The ladder, envelopes, gates, identities, action controls, records, and owners are documented and consistently applied.
4 - Managed	Delegated Autonomy performance, exceptions, incidents, value, downgrade, recovery, and renewal are measured and connected to decisions.
5 - Optimized	The organization continuously reallocates human and machine authority based on evidence while improving resilience, value, trust, and learning.

13 First Moves

Begin with one consequential but bounded workflow

1. Choose one workflow where employees already use AI or automation and where an action can affect money, records, customers, employees, or operations.
2. Map the workflow step by step and assign the current delegated autonomy level to each step.
3. Name the accountable owner and identify which decisions must remain human.
4. Complete an Delegated Autonomy Envelope for the highest-authority step.
5. Define the baseline, acceptance thresholds, stop conditions, and recovery test before release.
6. Run the appropriate Evidence Gate, record the decision, and set an automatic expiry.
7. Review the runtime evidence and be willing to tighten, downgrade, or stop as readily as expand.

Monday-morning test
A useful first result is not “more delegated autonomy.” It is visible authority, better evidence, safer action, and a decision the accountable executive can explain.

14 Source Notes

Evidence and internal TIO foundations

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 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
E7	NIST SP 800-160 Vol. 2 Rev. 1, Developing Cyber-Resilient Systems	https://csrc.nist.gov/pubs/sp/800/160/v2/r1/final
E8	Model Context Protocol, Security Best Practices	https://modelcontextprotocol.io/docs/tutorials/security/security_best_practices
E9	Regulation (EU) 2024/1689, Artificial Intelligence Act	https://eur-lex.europa.eu/eli/reg/2024/1689/oj
E10	Canada, Personal Information Protection and Electronic Documents Act	https://laws-lois.justice.gc.ca/eng/acts/P-8.6/
T1	TIO End-State Architecture Blueprint, Version 1.2 FINAL	Internal TIO Library
T2	TIO Sovereign Enterprise Intelligence Blueprint, Version 1.2 FINAL	Internal TIO Library
T3	TIO Factory Domain Guide, Version 1.0 FINAL	Internal TIO Library
T4	TIO Leadership Domain Guide, Version 1.0 FINAL	Internal TIO Library