Capability Guide: Governance

Speed Needs Decision Rights

Document Version 1.0 FINAL | TIO-KR-2026.08.0 | August 2026

Table of Contents

How to use this guide

Read Sections 1 through 3 for the executive view. Use Sections 4 through 7 to design and mature the capability. Use Section 8 to decide how to begin and where TSLA can help.

01 Executive Brief

The capability in one page

Governance is the organization's ability to set and deploy decision rights, accountabilities, policies, approval paths, and review rhythms for AI, automation, data, technology, risk, and value.

Good governance helps people move without guessing. When rules are unclear, approvals are hidden, or risk acceptance is informal, work either stalls, hides, duplicates, or moves without trust.

Element	Governance in plain English
Capability ID	L3 - Governance (Leadership domain)
Definition	The organization's ability to set decision rights, accountabilities, policies, approval paths, and review rhythms for AI, automation, data, technology, risk, and value.
Plain-English test	Do people know what is allowed, which policy applies, who decides, how to get approval or an exception, who owns the outcome, and where the decision is recorded?
Why it matters	AI and automation create speed. Speed without usable policy and clear decision rights creates shadow use, hidden risk, duplicated effort, and slow escalation.
Three outputs	Governance Model and Charter; AI Policy Suite and Approval Path; Decision, Exception, and Incident Record.
Accountable owner	CEO, owner, executive sponsor, transformation sponsor, or governance lead, depending on scale.
Maturity Level 0	AI and automation use is mostly invisible. Decisions depend on individuals, informal approvals, vendors, or avoidance.
Maturity Level 5	Governance is a source of speed, trust, learning, external confidence, and accountable innovation.
First move	Make current AI and automation use visible, then publish a minimum viable policy and one obvious approval path.
Executive point: Good governance answers six questions: 1. What is allowed? 2. Who can say yes? 3. Who can say no? 4. Who must be consulted? 5. Who accepts risk? 6. What evidence is required before work moves forward?	

2026.08 ALIGNMENT | Governance in the refreshed TIO system

Primary outcome contribution: Enterprise Autonomy

- Define decision rights for Enterprise Autonomy and Delegated Autonomy, including approved tool lanes, identities, authority limits, Evidence Gates, stop conditions, and residual-risk acceptance.
- Maintain an intelligent-system and agent register and apply the Assurance and Operations lifecycle.
- Require proportionate affected-stakeholder impact, challenge, appeal, and recourse for material or rights-affecting work.

Evidence at managed maturity: Governance is defined, owned, applied consistently, measured in operation, reviewed through evidence, and improved without weakening trust, accountability, or Enterprise Autonomy.

First refresh move: List the active AI tools and agents, name owners, and define the first approved use and authority lanes.

02 The Problem This Capability Solves

Why unclear decisions slow or expose the organization

Most organizations do not lack AI ideas. They lack a clear and usable way to decide what should happen with those ideas.

Employees are experimenting. Vendors are adding AI into existing systems. Teams are using personal subscriptions. Technology, Risk, Legal, Privacy, Security, Data, and People teams are often brought in after enthusiasm has already created commitment. Managers hear that AI matters but cannot always explain what is allowed.

Two gaps appear repeatedly. Some organizations publish an AI policy but do not define the operating system behind it: owners, approvals, exceptions, evidence, and review. Others create committees and controls but never translate them into rules employees and managers can understand. Governance must connect both sides.

The governance rule
Use the lightest structure that people will follow and leaders can trust. Governance should create visible use, plain policy, fast routing, recorded decisions, and learning that improves the system.

The visibility rule: You cannot govern what you cannot see. Governance starts with visibility, not control.

03 What this Capability Is

Clear boundaries, policy layers, and TIO relationships

Governance is the organization's ability to decide how decisions are made, then translate those decisions into policies and paths people can use. It tells people what is allowed, what needs review, who decides, where exceptions go, and how the organization learns from what happens.

Governance is	Governance is not
A decision system for AI, automation, data, technology, risk, and value.	A committee that reports status but rarely decides.
A plain-language policy layer connected to real approvals and controls.	A policy publication treated as proof that the work is governed.
A way to make current use, tools, and decisions visible.	The same as Risk, Legal, IT, Security, Privacy, or Compliance.
A clear approval, exception, and escalation path.	A substitute for Strategy or accountable business ownership.
A record of decisions, exceptions, incidents, and lessons learned.	A way to ban useful experimentation.
A rhythm for improving policies and decisions as the organization learns.	A one-time approval gate or annual policy refresh.

Where policy fits

Policy is the practical interface between leadership intent and daily work. It tells people what they must, may, and must not do. The governance system behind the policy supplies ownership, approvals, exceptions, evidence, and review.

Layer	What it does	Typical expression
Principles	State the commitments the organization intends to protect.	AI Guiding Principles or external trust commitments.
Policy	Set mandatory rules, responsibilities, permissions, and prohibitions.	AI Acceptable Use Policy.
Standards	Define specific requirements for tools, data, human review, records, vendors, and disclosure.	Approved tool list, data rules, review standards.
Procedures	Explain how to request approval, seek an exception, report an incident, or escalate a concern.	Approval path, exception route, incident process.
Evidence	Show that decisions are being made, policies are being followed, and the system is improving.	Decision records, training evidence, reviews, audits.

Keep this simple. These layers do not need five separate documents. A smaller organization may combine them into one short policy, one approved-tool page, and one approval and incident path.

In TIO, Governance sits in the Leadership domain. Leadership owns the design of decision rights, policy, accountability, and review. Data, Technology, Factory, and People deploy those rules inside their work. The Leadership logic remains simple: Strategy chooses. Fluency informs. Governance decides. Risk protects. Transformation mobilizes.

Boundary rule **Leadership designs the governance system. Risk protects trust. Data, Technology, Factory, and People apply the policies, standards, and decision rights in their own work.**	
Behavior	**Meaning**
See	Make current use, tools, use cases, data, decisions, exceptions, and incidents visible.
Guide	Give people plain-language principles, policies, standards, risk tiers, approved tool lanes, and examples.
Decide	Clarify who approves, rejects, escalates, funds, scales, pauses, stops, or accepts risk.
Control	Apply review thresholds, data rules, human oversight, evidence requirements, exceptions, and escalation paths.
Improve	Use decisions, policy questions, exceptions, incidents, near-misses, and value evidence to improve governance.
TIO area	**Governance relationship**
Strategy	Uses enterprise choices and risk posture to set priorities, decision thresholds, and policy boundaries.
Fluency	Gives leaders enough judgment to apply policy, challenge weak claims, and sponsor work responsibly.
Risk	Defines risk treatment and routes work to the right review and risk-acceptance owner.
Data	Applies policies for ownership, quality, access, use, retention, and evidence.
Technology	Implements approved, restricted, pilot-only, and prohibited tool lanes and technical controls.
Factory	Connects AI Use Case approval to Evidence Gates, value, scaling, redesign, and stop decisions.
People	Explains policies, supports managers, builds skills, and helps people raise concerns without fear.

04 How the Capability Works

Visibility, routing, decisions, evidence, and learning

Governance should work like a useful operating rhythm. A person with a question or idea should be able to see the policy, classify the work, find the right approval path, and receive a decision without guessing.

A simple Governance flow:

Current use or new idea -> visibility check -> context, data, and risk classification -> policy and tool-lane check -> approval or exception path -> decision and owner -> evidence record -> monitor -> improve policy and governance

The flow should be fast for green work, careful for yellow work, and deliberate for red work. The tier is determined by data sensitivity, potential impact, reversibility, and who could be affected.

AI is often where the need appears first, but the capability also governs automation, data, technology, value, and transformation decisions. The same rule applies: make the work visible, route it clearly, name the owner, record the decision, and learn from the outcome.

Use context	Typical question	Governance need
Personal use	Can an individual use AI to help with their own work?	A short acceptable-use policy, approved or restricted tools, a data boundary, and clear examples.
Team use	Can a team use AI in a recurring workflow?	Named owner, use-case entry, data boundary, human review, team guidance, and approval where required.
Enterprise use	Can this become part of an enterprise system or supported process?	Technology, data, security, vendor, support, change, value, and approval decisions.
External use	Can AI affect a customer, resident, patient, employee, partner, regulator, or the public?	Stronger review, disclosure where appropriate, human oversight, assurance evidence, incident response, and executive accountability.

The Governance Legibility Lens tests whether the system can carry clear intent, honest feedback, shared learning, and external trust. It uses four information flows and eight things the organization must make clear.

Direction	What must travel	Failure mode it prevents
Leadership to work	Strategy, risk posture, policy, tool lanes, decision rights, and boundaries reach daily work.	Policy fog: people hear "use AI" and "do not create risk" but cannot tell what to do.
Work to leadership	Real use, hidden workflows, near-misses, friction, value signals, and exceptions move upward.	Executive illusion: leaders believe the dashboard while work runs on shadow AI.
Team to team	Use cases, prompts, tools, incidents, vendor lessons, and good practice move across teams.	Fragmented learning: every team pays the learning cost alone.
Inside to outside	The organization can explain its AI posture, data handling, human oversight, and trust commitments.	Trust gap: customers and partners cannot see why the organization can be trusted.

What Governance must make clear	Plain-English meaning
Policy and principles	Rules and commitments people can read, explain, and use.
Use cases	A living catalog of what AI and automation are used for, by whom, and why.
Risk tiers	A green, yellow, red routing model people can apply before asking for approval.
Approval and exceptions	One obvious route for review, approval, exception, escalation, pause, or stop.
Incidents and escalation	A blameless way to report mistakes, near-misses, wrong outputs, data issues, and external impacts.
Tools and training	Approved, restricted, pilot-only, and prohibited tools, with practical guidance and training.
Communication and learning	Regular review, show-and-tell, feedback, and lessons that move across the organization.
External trust	Clear customer or public commitments on data, security, oversight, disclosure, assurance, and incident response.

Show-and-tell helps people learn and reduces hidden use. It does not replace approval. Sensitive data, repeatable workflows, external impact, material decisions, or enterprise deployment still need the appropriate route.

A simple monthly Governance review should examine new use cases, policy questions, exceptions, incidents, near-misses, tool lessons, external risks, decision speed, and rules that need to change.

Tier	Typical examples	Governance response
Green	Public information, personal productivity, internal drafts with no sensitive data, low-impact research.	Allowed within policy. Record only if it becomes a repeatable team or enterprise use case.
Yellow	Internal workflows, recurring team use, non-public information, operational analysis, limited external content.	Needs a named owner, data boundary, human review, and the defined approval path.
Red	Customer or employee data, financial, legal, regulated, cybersecurity, safety, high-impact decisions, or external-facing automation.	Needs formal review by accountable leaders and the relevant business, Risk, Legal, Privacy, Security, Data, Technology, and People owners.

05

Outputs, Evidence, and Measures

What Governance must produce and prove

Governance should produce a small number of useful outputs. Three are enough to make the capability real without creating a policy library or governance bureaucracy.

Output	Purpose	Quality standard
Governance Model and Charter	Defines ownership, decision rights, review scope, risk acceptance, cadence, escalation, evidence expectations, and improvement rhythm.	Short enough to use, specific enough to decide, and clear enough that leaders can explain it.
AI Policy Suite and Approval Path	Combines the internal acceptable-use policy, approved-tool and data rules, approval and exception route, and external guiding principles where needed.	Plain English, proportionate to risk, connected to real controls, and maintained by a named owner.
Decision, Exception, and Incident Record	Creates the evidence trail for approvals, rejections, exceptions, incidents, near-misses, residual-risk decisions, and Evidence Gate outcomes.	Used in governance reviews. The record must improve decisions, not become an archive.
Evidence type	**Examples**	
Existence evidence	Approved charter; acceptable-use policy; external principles where needed; tool and data standards; risk tiers; approval and exception path; decision and incident record.	
Use evidence	Employees and managers know the rules; use cases are routed; decisions and exceptions are recorded; incidents are reported; tool lanes are followed; policies are used in real work.	
Value evidence	Faster and clearer decisions; fewer hidden tools; fewer avoidable incidents; better use-case quality; clearer risk acceptance; better scaling decisions; stronger external trust.	
Measure	**What it tells us**	**Caution**
Policy clarity	Whether employees and managers know what is allowed and where to ask.	A published policy does not prove understanding.
Decision-cycle time	Whether the approval path helps work move.	Do not reward speed that bypasses material review.
Governance coverage	Whether active tools and material use cases have owners, classifications, and current rules.	Coverage should focus on material exposure, not paperwork volume.
Exception pattern	Where policy is unclear, too strict, too loose, or out of date.	Exceptions should improve the system, not punish candour.
Incident and near-miss reporting	Whether people trust the reporting channel enough to surface problems.	Low reporting may mean hidden issues, not low risk.

Output	Purpose	Quality standard
Evidence Gate outcomes	Whether approved work creates value and should scale, change, or stop.	Do not let every pilot graduate by default.

Assessment crosswalk. Use this guide with the TIO Capability Model and Maturity Assessment Tool. Scores should be supported by evidence that governance exists, is understood, is used, and improves decisions.

Assessment field	Governance answer
Capability	Governance, Leadership domain.
Current maturity	Score using the 0 to 5 maturity table in Section 6.
Target maturity	Set based on operating reality, strategic importance, customer exposure, regulatory pressure, and current pain.
Evidence reviewed	Charter, policy suite, approved-tool and data rules, approval path, use survey, use-case catalog, decision and exception record, incident route, governance agendas, and manager or employee clarity evidence.
Evidence confidence	High only when policies and decision paths are used in real work, not merely drafted or published.
First move	Run a visibility scan, publish a minimum viable acceptable-use policy, define the approval and exception path, and start the decision record.
Owner	Executive sponsor plus named Governance owner.
Priority logic	Strategic importance x current pain x maturity gap / effort.

06

Maturity: 0 to 5

How to score the capability honestly

Use the maturity model to score the capability honestly. A policy proves that a rule exists. Maturity requires evidence that people understand it, decisions follow it, exceptions and incidents are learned from, and the system improves.

Level	What it looks like	Evidence	Next move
0 Absent	No meaningful Governance capability exists. AI and automation use is invisible or avoided. Decisions depend on individuals, vendors, or informal permission.	No clear policy, approval path, catalog, decision record, or incident route.	Name the issue and accountable owner. Make current use visible.
1 Initial / Ad-Hoc	A draft policy or scattered rules and approvals exist, but they are uneven and person-dependent.	Draft policy, informal tool guidance, isolated approvals, local use cases.	Create repeatable basics: minimum policy, risk tiers, named approvers, approval path, and record.
2 Repeatable	Approved plain-language policy and basic governance exist for common activity. People can route work more than once.	Acceptable-use policy, green/yellow/red routing, approved tools, basic approval and exception path, early decision record.	Document the model, explain it to managers, and use it consistently.
3 Defined	Governance is documented, owned, communicated, and used consistently for priority work.	Governance Charter, policy suite, use-case catalog, tool lanes, incident route, manager guidance, regular review cadence.	Measure policy clarity, decision speed, exceptions, incidents, and Evidence Gate outcomes.
4 Managed	Governance is measured, reviewed, and connected to strategy, risk, portfolio, policy management, and value decisions.	Clarity results, coverage, cycle time, exception and incident trends, policy updates, value reviews, improvement actions.	Use evidence to simplify, strengthen, and update the system.
5 Optimized	Governance continuously adapts and helps the organization move faster with trust. Decision rights are federated where appropriate.	Adaptive policy and standards, learning loops, trusted external posture, federated rights, continuous improvement record.	Keep the system light, current, evidence-based, and connected to real decisions.
How to score honestly A policy alone proves existence, not maturity. A committee alone proves structure, not effectiveness. A mature score requires real use cases to be routed, real decisions and exceptions to be recorded, and policies and decision paths to improve based on evidence.			

07 How to Build and Mature the Capability

A practical path from visibility to adaptive Governance

Start light. Build only what helps people see, understand, route, decide, record, and learn. Add more policy and control only when risk, scale, customer exposure, regulation, or complexity requires it.

Minimum viable Governance capability:

- One executive sponsor and one accountable Governance owner.
- One current-use survey or discovery conversation.
- One short internal AI Acceptable Use Policy.
- One approved, restricted, pilot-only, and prohibited tool view, with a green/yellow/red routing model.
- One obvious approval and exception path.
- One lightweight decision, exception, and incident record.
- One blameless incident and near-miss channel.
- One monthly review rhythm during active build.

Stage	What to build	What to avoid
First 30 days	Name the sponsor and owner. Run a use scan. Draft a minimum acceptable-use policy. Create temporary risk tiers, approval and exception path, and decision record. Review the top current use cases.	Do not begin with a large policy project or another tool purchase.
First 90 days	Approve the Governance Charter and policy suite. Publish the tool and data rules. Brief managers. Build the use-case catalog. Open the incident channel. Connect decisions to Factory Evidence Gates.	Do not let policy sit apart from tools, workflows, approvals, and manager behaviour.
First 12 months	Measure policy clarity, decision speed, governance coverage, exceptions, incidents, use-case outcomes, and value evidence. Update policies as the organization learns. Strengthen external trust commitments where needed.	Do not confuse policy volume or review volume with governance quality.
Year 2+	Keep decision rights, policies, standards, and evidence current as tools, laws, vendors, business models, risks, and stakeholder expectations change.	Do not lock Governance into yesterday's technology assumptions.
Policy lifecycle	Draft from real use and risk. Approve through named decision rights. Explain in plain language. Apply through tools, approvals, and controls. Improve through questions, exceptions, incidents, evidence, and scheduled review. Do not publish and forget.	

Maturity path

Stage	Governance focus
Crawl	Make current use visible. Set a safe-use policy. Create one approval and exception path.
Walk	Build the use-case catalog, risk tiers, tool lanes, manager guidance, incident channel, and decision record.
Run	Connect policy and Governance to the transformation portfolio, Evidence Gates, external-facing use, and value review.
Fly	Governance becomes adaptive, evidence-based, federated where needed, and trusted inside and outside the organization.

Scale-fit application. The capability remains the same. The amount of policy, formal review, documentation, and assurance changes with operating reality.

Operating reality	What Governance should look like
Owner-led business	One-page acceptable-use policy, owner-led decision rights, one approved-tool view, one approval and incident route, and direct review of anything sensitive or external-facing.
Managed enterprise	CEO or GM sponsor, short policy suite, small cross-functional review group, basic tool lanes, manager-carried routing, and monthly review.
Mid-sized enterprise	Formal but light Charter, internal policy, external principles where needed, cross-functional review, use-case catalog, risk tiers, decision record, incident route, and quarterly value review.
Large enterprise	Formal decision bodies, policy management, delegated rights, portfolio integration, model and tool review, legal/privacy/security controls, audit evidence, and board visibility.
Large distributed enterprise	Enterprise policies and standards with local decision rights, federated governance, shared records, local escalation, unit-level evidence, and enterprise learning rhythm.
Customer-facing step change Governance becomes more demanding when AI or automation affects customers, citizens, patients, employees, partners, regulators, or the public. Internal policy protects daily work. External principles explain the organization's commitments.	

Before customer-facing AI use scales, leaders should be able to answer: What promise does this support? What data is used? Where is human review required? What happens when output is wrong? What will be disclosed? Who owns the outcome? What incident route exists? What evidence supports continuation or scale?

08 TSLA Support and First Moves

How leaders can begin without building bureaucracy

TSLA helps leaders build Governance as a practical enterprise capability. The work connects decision rights, policy, approval paths, risk, evidence, and review so executives can govern and employees can act.

TSLA support area	What TSLA helps with
Executive Governance Briefing	Help leaders understand Governance as a decision system, policy layer, trust mechanism, and source of speed.
Governance Design Workshop	Define scope, principles, decision rights, risk acceptance, policy architecture, operating rhythm, escalation, and evidence expectations.
Governance Model and Charter	Develop a practical Charter that is short, clear, and connected to the TIO transformation system.
AI Policy Suite	Develop the internal Acceptable Use Policy, approved-tool and data rules, approval and exception path, and external Guiding Principles where needed.
AI Use Visibility Scan	Survey or interview leaders and teams to identify current tools, personal accounts, use cases, value, uncertainty, risks, and hidden workflows.
Decision, Exception, and Incident Record	Design the evidence record for approvals, rejections, exceptions, incidents, near-misses, risk acceptance, and Evidence Gate outcomes.
Governance Sprint	Build minimum viable Governance in 30 to 90 days without creating unnecessary bureaucracy.
Governance Rhythm Advisory	Support early reviews, escalation decisions, policy improvements, and quarterly learning and value rhythms.
External Trust Posture	Help organizations explain AI use, data handling, security, human oversight, disclosure, assurance, customer commitments, and incident response.
TSLA position Here is what the organization must own. Here is where TSLA can help build it faster, safer, and more practically.	

What to do Monday morning:

1. Name the executive sponsor and accountable Governance owner.
2. Ask what AI and automation tools are already being used, including personal accounts and embedded vendor features.
3. Draft one short internal acceptable-use policy: what is allowed, what is prohibited, and what needs review.
4. Publish the approved-tool and data boundary, with a simple green/yellow/red routing model.
5. Set up one approval and exception path and one decision record.
6. Open a blameless incident and near-miss channel.
7. Review the first five material use cases against strategy, value, data, policy, risk, human review, external impact, and ownership.

Final word
Good Governance gives people clear rules and leaders defensible decisions. Policy makes the rules usable. Decision rights, evidence, and review make them real.

09

Source Notes and Research Base

The research base and alignment behind this guide

This guide is aligned with the TIO Framework, the TIO Capability Model and Maturity Assessment Tool, the Strategy and Fluency Capability Guides, the TIO Implementation Guides, and client-tested Governance work. Client-specific details have been generalized.

External practice anchors include the NIST AI Risk Management Framework, ISO/IEC 42001, OECD AI Principles, risk-based regulatory approaches, enterprise risk practices, and current AI governance practice. They support accountable ownership, plain policy, risk-based routing, documentation, human oversight, transparency, monitoring, incident learning, and continuous improvement.

The public TIO rule remains simple: use the lightest structure that creates clarity, accountability, trust, and value. Simplicity is the ultimate sophistication.