Capability Guide: Risk

Trust Is a Leadership Responsibility

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

Risk is the organization's ability to identify, assess, control, monitor, escalate, and learn from the risks created by AI, automation, data, technology, and transformation work.

Good Risk does not stop progress. Weak Risk either blocks everything too late or lets the organization move without enough protection. Both create avoidable harm.

Element	Risk in plain English
Capability ID	L4 - Risk (Leadership domain)
Definition	The organization's ability to understand what could go wrong, how serious it could be, who could be affected, what protection is needed, and what residual risk leaders are accepting.
Plain-English test	Can leaders explain the main risks, the required controls, the accountable owner, the escalation path, and the residual risk before work moves forward?
Why it matters	AI, automation, data, technology, vendors, redesigned work, and changed roles can create trust, privacy, security, legal, operational, financial, reputational, and workforce exposure.
Five behaviors	Identify. Assess. Protect. Escalate. Learn.
Three practical outputs	Risk Profile and Trust Boundaries; Risk Assessment and Control Plan; Residual Risk and Incident Learning Log.
Accountable owner	Executive sponsor with support from Risk, Legal, Privacy, Security, Technology, Data, People, and business owners.
Maturity headline	At Level 0, risk is invisible or discovered late. At Level 5, risk learning helps the organization move faster because trust protection is built into how work changes.
First move	Name the Risk owner, define the first trust boundaries, and assess the first five AI or automation use cases using a simple risk lens.

Leadership rule:

Risk is not a specialist problem that leaders delegate away. Leaders own the trust posture of the organization. Specialists help them see, assess, control, and learn.

2026.08 ALIGNMENT | Risk in the refreshed TIO system

Primary outcome contribution: Enterprise Autonomy

- Assess model, agent, data, dependency, concentration, supply-chain, stakeholder-impact, skill-atrophy, operational, and recovery risks.
- Risk tiering should create fast, visible paths for low-risk uses and deeper review for material uses.
- Residual risk remains a named human decision supported by evidence and specialist input.

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

First refresh move: Tier current AI and automation uses by consequence, data sensitivity, reversibility, and affected stakeholders.

02 The Problem This Capability Solves

Why unprotected ambition damages trust

Most organizations are moving faster than their risk systems. People are trying AI tools. Vendors are embedding AI into existing platforms. Teams are automating work. Data is moving across systems. Leaders are approving pilots before the risk picture is clear.

The problem is not ambition. The problem is unprotected ambition.

When Risk is weak, the organization discovers exposure late. A tool has already been used. A workflow has already changed. A vendor has already been selected. Customer data has already moved. An employee impact has already appeared. A public or client-facing claim has already been made.

That is expensive. It slows decisions, creates rework, damages confidence, and can put trust at risk.

When Risk is weak	What happens
Risk is invisible	Leaders do not know where AI, automation, data, or vendor exposure already exists.
Risk is late	Problems are found after decisions have been made or after work has already scaled.
Risk is vague	People know something feels risky, but cannot explain what could go wrong or what control is needed.
Risk is overdone	Everything is treated as high risk, so useful work stalls.
Risk is underdone	Low-friction experiments quietly become recurring processes without review.
Risk is disconnected	Privacy, security, legal, technology, people, operations, and business owners see pieces, but no one owns the whole risk picture.

The trust problem:

Risk protects trust before trust is damaged. That includes customer trust, employee trust, board trust, regulator trust, partner trust, and public trust.

03 What this Capability Is

Clear boundaries, trust protection, and TIO relationships

Risk is the capability that helps the organization ask better questions before work moves forward and learn faster when something goes wrong.

In TIO, Risk sits in the Leadership domain because risk posture is a leadership choice. The Leadership logic is simple: Strategy chooses. Fluency informs. Governance decides. Risk protects. Transformation mobilizes.

Risk is broader than AI risk. AI is often where the need shows up first, but this capability also covers automation, data, technology, vendors, operating changes, workforce impact, and transformation risk.

Risk is	Risk is not
A practical way to see what could go wrong before work scales.	The department that says no to everything.
A trust-protection capability for AI, automation, data, technology, and transformation work.	A substitute for Governance, Legal, Security, Privacy, Technology, Data, Factory, or People.
A way to define controls, oversight, escalation, and residual risk.	A checklist completed after the real decision has already been made.
A learning system that improves from incidents, near-misses, exceptions, and evidence.	A binder of risk language that does not change behavior.
A leadership discipline supported by specialists.	A specialist function leaders can ignore or outsource entirely.

Boundary rule

Governance decides who can approve, stop, escalate, and review. Risk identifies what could go wrong, how serious it is, what controls are needed, and whether trust is protected.

Behavior	Meaning
Identify	Find where AI, automation, data, technology, vendor, operational, workforce, customer, or public exposure may exist.
Assess	Determine severity, likelihood, affected parties, data sensitivity, legal or regulatory exposure, and business impact.
Protect	Define the controls, safeguards, human oversight, testing, monitoring, and communication needed to reduce risk.
Escalate	Route serious risks and residual-risk decisions to the right accountable leaders.
Learn	Use incidents, near-misses, exceptions, audits, and Evidence Gate outcomes to improve future rules and controls.

04 How the Capability Works

Risk lenses, tiers, controls, escalation, and learning

Risk should be built into the path from idea to decision. It should not arrive as a late-stage objection.

A simple Risk flow

New idea or existing use -> risk lens -> risk tier -> control plan -> residual risk decision -> monitoring -> learning loop

The work starts with three plain questions: What could go wrong? Who could be affected? What protection is needed before this moves forward?

Risk lens	What to ask
Trust	Could this damage customer, employee, partner, board, regulator, or public confidence?
Data	What data is used? Is it public, internal, confidential, personal, regulated, customer, employee, financial, or commercially sensitive?
Security	Could the work create a new attack surface, access risk, leakage risk, or dependency on an unapproved tool or vendor?
Legal and compliance	Could the work trigger legal, contractual, regulatory, procurement, records, privacy, accessibility, or audit obligations?
Operations	Could the work disrupt service, create process errors, reduce resilience, or make the organization dependent on fragile automation?
Financial and value	Could the work create hidden cost, weak return, bad incentives, or value claims the organization cannot prove?
People	Could the work affect roles, work quality, fairness, safety, skills, job security, accountability, or manager behavior?
External use	Could output, insight, advice, or automation reach a customer, citizen, supplier, regulator, media outlet, or public audience?

The same use case can move from low to high risk when the context changes. A draft using public information may be low risk. A recurring workflow using customer data, employee data, legal content, financial decisions, or external-facing output needs stronger review.

Risk tier	Typical examples	Risk response
Green	Public information, individual productivity, drafting with no sensitive data, exploratory learning.	Allowed within safe-use rules. No formal risk review unless it becomes recurring, shared, automated, or external-facing.
Yellow	Internal workflows, team use, non-public data, recurring reporting, operational analytics, manager-facing decision support.	Needs owner, risk assessment, data boundary, human review, control plan, and approval path.
Red	Customer or employee data, legal, financial, compliance, cybersecurity, regulated, automated decisions, external-facing output, or material workforce impact.	Needs formal review by accountable leaders and relevant Risk, Legal, Privacy, Security, Data, Technology, People, and business owners.

Risk does not own every control. It makes sure the right control exists, has an owner, is proportionate, and is reviewed when evidence changes.

Decision point	What Risk contributes
Before approval	Risk tier, exposure summary, required controls, review partners, residual-risk question.
Before pilot	Testing needs, data boundary, human review point, monitoring expectation, incident route.
Before scale	Evidence of control performance, unresolved risks, exceptions, user impact, customer impact, value proof.
After incident or near-miss	Root cause, corrective action, affected parties, control change, communication need, rule update.

05

Outputs, Evidence, and Measures

What Risk must produce and prove

Risk should produce a small number of useful outputs. The goal is not more risk paperwork. The goal is better decisions and stronger trust protection.

Output	Purpose	Quality standard
Risk Profile and Trust Boundaries	Names the main risk categories, the organization's trust boundaries, and the areas where leaders require stronger review.	Clear enough for leaders and managers to use. Specific enough to guide real decisions. Updated as tools, work, customers, regulation, and evidence change.
Risk Assessment and Control Plan	Assesses a use case, tool, workflow, vendor, automation, or transformation change and defines the controls needed before it moves forward.	Proportionate to the risk tier. Practical controls. Named owners. Clear human review, monitoring, and escalation points.
Residual Risk and Incident Learning Log	Records accepted residual risk, exceptions, incidents, near-misses, corrective actions, lessons learned, and control changes.	Used in governance and risk reviews. If the log does not improve decisions, it is only an archive.
Evidence type	**Examples**	**What it proves**
Existence evidence	Risk profile, trust boundaries, risk categories, risk-tier model, assessment template, control library, incident route.	The capability has been designed.
Use evidence	Completed risk assessments, control plans, residual-risk approvals, incident reviews, near-miss records, updated rules.	The capability is being used in real decisions.

Output	Purpose	Quality standard
Value evidence	Fewer late surprises, faster risk reviews, fewer unresolved exceptions, stronger customer assurance, improved audit readiness, better Evidence Gate decisions.	The capability is protecting trust and helping work move safely.

Measure	What to track
Use-case risk coverage	Percentage of yellow and red AI or automation use cases with a completed risk assessment.
Control readiness	Percentage of material use cases with named controls, owners, and review points before pilot or scale.
Residual-risk discipline	Number of residual-risk decisions recorded with accountable owner and rationale.
Incident learning	Number of incidents and near-misses reviewed, closed, and used to improve controls or rules.
Decision speed	Average time from risk intake to risk recommendation by tier.
External trust readiness	Percentage of customer-facing uses with data handling, human oversight, escalation, and assurance evidence documented.

Assessment crosswalk. Use this guide with the TIO Capability Model and Maturity Assessment Tool. Scores should be supported by specific evidence, not aspiration.

Assessment field	Risk answer
Capability	Risk, Leadership domain.
Current maturity	Score using the 0 to 5 maturity table in Section 6.
Target maturity	Set based on strategic importance, customer exposure, data sensitivity, regulatory environment, operating complexity, and transformation pace.
Evidence reviewed	Risk profile, trust boundaries, assessments, control plans, residual-risk approvals, incident and near-miss log, corrective actions, governance records.
Evidence confidence	High only when risk evidence is used in decisions, not merely stored.
Owner	Executive sponsor with accountable Risk owner and support from relevant specialist functions.
First move	Assess the first five AI or automation use cases and define the minimum trust boundaries.
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 Risk honestly. A high score requires evidence that Risk exists, is used, and improves protection, decisions, and trust.

Level	What it looks like	Evidence	Next move
0 Absent	No meaningful Risk capability exists for AI, automation, data, technology, or transformation work. Risk is invisible or discovered after harm has occurred.	No risk lens, no trust boundaries, no assessments, no control plan, no residual-risk record, no incident learning loop.	Name the exposure, name an owner, and assess the first visible use cases.
1 Ad Hoc	Some risks are discussed, but review depends on individuals, relationships, or late escalation.	Scattered reviews, informal advice, inconsistent approvals, isolated risk notes.	Create repeatable basics: risk lens, green/yellow/red routing, assessment template, and incident route.
2 Repeatable	Basic risk assessment exists for common use cases. People have a way to identify risk, define controls, and escalate when needed.	Risk-tier model, basic assessments, simple control plans, named reviewers, early incident log.	Document the operating model, communicate it, and connect it to Governance and Evidence Gates.
3 Defined	Risk is documented, owned, communicated, and used consistently for priority AI, automation, data, technology, and transformation work.	Risk profile, trust boundaries, assessment records, control plans, residual-risk approvals, incident learning records.	Measure risk review speed, control performance, exceptions, incidents, and learning.
4 Managed	Risk is measured, governed, improved, and connected to strategy, portfolio, vendor, data, technology, people, and value decisions.	Dashboards, risk trends, control effectiveness evidence, open risk items, post-incident reviews, audit-ready records.	Use risk evidence to improve investment choices, governance rules, controls, and transformation priorities.

Level	What it looks like	Evidence	Next move
5 Optimized	Risk continuously improves as tools, regulations, customers, workforce impact, and evidence change. It helps the organization move faster because trust protection is built in.	Adaptive controls, fast risk routing, learning loops, trusted customer posture, mature residual-risk discipline.	Keep simplifying. Retire controls that no longer help and strengthen controls where evidence shows exposure.

07 How to Build and Mature the Capability

A practical path from visible exposure to adaptive protection

Start light. Build only what helps people identify, assess, protect, escalate, and learn. Add more structure only when risk, scale, customer exposure, regulation, or complexity requires it.

Minimum viable Risk capability:

- One executive sponsor and one accountable Risk owner.
- One simple risk lens for AI, automation, data, technology, and transformation work.
- One set of trust boundaries that names what needs stronger review.
- One green/yellow/red risk-tier model.
- One risk assessment and control-plan template.
- One residual-risk approval route.
- One incident and near-miss learning log.
- One monthly review rhythm during active build.

Stage	What to build	What to avoid
First 30 days	Name sponsor and owner. Define the first trust boundaries. Create a simple risk lens. Assess the first five visible AI or automation use cases. Open an incident and near-miss route.	Do not start with a large enterprise risk manual or a tool purchase.
First 90 days	Finalize risk categories, routing, assessment template, control-plan format, residual-risk approval path, and reporting rhythm. Connect Risk to Governance and Evidence Gates.	Do not let every review become bespoke. Repeatability matters.
First 12 months	Measure review speed, control readiness, residual-risk decisions, incidents, near-misses, corrective actions, and customer-facing assurance evidence.	Do not confuse risk documentation with trust protection.
Year 2+	Refine controls based on evidence. Integrate risk learning with strategy, governance, technology, data, people, vendor, and Factory decisions.	Do not lock Risk into yesterday's tools, policies, or threat assumptions.

Scale-fit application. The capability stays the same. The machinery changes by operating reality.

Operating reality	What Risk should look like
Owner-led business	Owner-led trust boundaries, simple safe-use rules, visible tools and use cases, one risk lens, and a small number of must-review areas.
Managed enterprise	CEO or GM sponsor, named Risk owner, practical risk categories, simple assessment process, control plans for priority use cases, and leadership review.
Mid-sized enterprise	Formal but light risk profile, cross-functional review, risk-tier routing, control plans, residual-risk log, and customer-facing trust posture where needed.
Large enterprise	Integrated risk governance, policy management, model/tool/vendor risk review, control libraries, dashboards, audit support, and portfolio-level reporting.
Large distributed enterprise	Federated risk model, common standards, local accountability, enterprise reporting, stronger regulatory alignment, and consistent incident learning across units.

08

TSLA Support and First Moves

How leaders can begin without creating risk bureaucracy

TSLA helps leaders build Risk as a practical trust-protection capability, not as a paperwork exercise. The work is designed to be clear enough for executives, useful enough for managers, and strong enough to support safe movement.

TSLA support area	What TSLA helps with
Executive Risk Briefing	Educate leadership on Risk as trust protection, not late-stage objection.
Risk Design Workshop	Define risk categories, trust boundaries, review partners, escalation, residual-risk discipline, and learning rhythm.
Risk Profile and Trust Boundaries	Develop a practical risk profile that fits the organization's strategy, operating reality, customer exposure, and maturity.
Risk Assessment and Control Plan	Create assessment templates, control-plan formats, and risk-tier routing that teams can actually use.
Residual Risk and Incident Learning Log	Build a simple evidence trail for accepted risk, exceptions, incidents, near-misses, corrective actions, and lessons learned.
Governance and Risk Integration	Connect Risk to Governance decision rights, AI Use Case intake, Evidence Gates, and transformation reviews.
External Trust Posture	Prepare customer, board, regulator, or partner-facing language where AI or automation affects external trust.
Thought-Partnership	Support leaders as they make difficult risk, trust, and transformation decisions.

What to do Monday morning:

1. Name the executive sponsor and accountable Risk owner.
2. List the first five AI, automation, data, technology, or transformation uses that could create trust exposure.
3. Define the first trust boundaries: what is low risk, what needs review, and what must not proceed without approval.
4. Create a simple risk lens using trust, data, security, legal, operations, financial, people, and external-use questions.
5. Assess the first five use cases and assign a green, yellow, or red risk tier.
6. For yellow and red work, name the required controls, owners, human review points, and escalation path.
7. Start a residual-risk and incident learning log.

Final word

Risk is not about making the organization less ambitious. It is about making ambition safer, clearer, and more trusted. In the age of AI and AI and automation, trust is not protected by caution alone. Trust is protected by disciplined leadership.

09

Source Notes and Research Base

The research base and alignment behind this guide

This guide is part of the TIO Library. It is aligned to the TIO Framework, the TIO Capability Model and Maturity Assessment Tool, the finalized Governance Capability Guide structure, and the working TIO Leadership logic: Strategy chooses. Fluency informs. Governance decides. Risk protects. Transformation mobilizes.

It is written as a practical executive guide. It is not legal, regulatory, privacy, security, financial, or professional risk advice. Organizations should apply their own professional review where required.