Data Trust & Reliability

A Practical Authority Framework for Decision-Critical and AI-Driven Organizations

Trust Your Data. Prove It.

A Practical Authority Framework for Decision-Critical and AI-Driven Organizations

Martijn Wiggers

Independent Data Trust Authority & Founder of FactVault

In collaboration with

Komawi Development

www.komawi.com

Copyright & Usage

This document and its contents are © 2026 Martijn Wiggers.

Permission is granted to share this document unmodified for educational and internal organizational use, provided attribution is retained.

Commercial redistribution, modification, or incorporation into proprietary frameworks requires explicit permission.

This framework is intentionally normative.
Its structure, terminology, and reliability scale form an integrated whole.
Selective adoption without preserving definitions and boundaries may invalidate its conclusions.

This document is maintained at https://factvault2.komawi.com

Executive Summary

Data Trust as a Strategic Capability for AI and Decision-Making

The problem

Organizations increasingly rely on data and AI to drive decisions, automation, and strategy. While investments in analytics, data platforms, and AI models continue to grow, many organizations struggle with a more fundamental issue:

They cannot clearly explain why their data can be trusted for the decisions it is used to make.

Data quality metrics, model confidence scores, and sophisticated pipelines do not solve this problem. They often mask it.

The result is a widening gap between:

- what organizations *can* automate, and
- what they can *defend* under scrutiny.

This gap is now the primary constraint on scalable, responsible AI adoption.

The core insight

Data trust is not a technical attribute. It is an organizational capability.

Trustworthy data use requires explicit answers to questions such as:

- Where did this data come from?
- Under which conditions is it valid?
- Who is accountable for its correctness?
- What decisions are we allowed to make based on it?
- When must automation stop and human judgment intervene?

Most organizations answer these questions implicitly.
AI systems, however, require explicit answers.

Reliability vs quality, accuracy, and confidence

A central distinction in this framework is between four commonly conflated concepts:

- **Accuracy** — Is the value correct?
- **Quality** — Is the data fit for a specific purpose?
- **Confidence** — How certain is a model or human about the result?
- **Reliability** — *May we responsibly act on this data?*

Only reliability determines whether data may be used for decisions that are automated, customer-facing, regulatory, or irreversible.

High accuracy or confidence does not imply reliability.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Trust decay is inevitable — unless managed

As data moves through systems, it is copied, transformed, aggregated, enriched, and inferred. Each step risks:

- loss of provenance,
- loss of context,
- and dilution of accountability.

This phenomenon—**trust decay**—often occurs silently, even when data remains technically correct.

Without explicit controls, organizations accumulate data that looks authoritative but is no longer defensible.

Context multiplies reliability

Facts without context are incomplete truths.

Context—legal, contractual, temporal, or situational—determines how far trust extends. It can:

- increase reliability by clarifying meaning, or
- constrain reliability by exposing exceptional conditions.

Context is fragile and easily lost. Once lost, it cannot reliably be reconstructed.

Trustworthy systems treat context as first-class information.

Decision boundaries make trust operational

Trust should not be treated as a score to optimize.
It should function as a **gate** that limits action.

Different reliability levels permit different types of decisions:

- exploratory analysis,
- advisory recommendations,
- human-approved actions,
- or fully automated execution.

Without explicit decision boundaries, AI systems will inevitably exceed what trust can justify.

Human intervention is a design principle

Human involvement is not a failure mode of automation.
It is a **necessary control mechanism**.

Human judgment is required at defined control points, including:

- source onboarding,
- context validation,
- trust escalation,
- exception handling,
- and corrections.

Human-in-the-loop works only when it is deliberate, bounded, and accountable.

© 2026 Martijn Wiggers
This document provides conceptual guidance.
It does not constitute legal or regulatory advice.

All rights reserved.

Corrections do not automatically restore trust

Fixing a value does not necessarily restore reliability.

Trust recovery depends on:

- availability of evidence,
- preservation of context,
- and continuity of accountability.

In some cases, permanent trust degradation is the correct and honest outcome.

Organizational readiness determines AI impact

Organizations ready to operate with trusted data exhibit:

- clear ownership and accountability,
- explicit decision responsibility,
- willingness to constrain automation,
- tolerance for uncertainty,
- and cross-functional alignment.

Technology alone cannot compensate for missing organizational readiness.

Why this matters now

AI is moving rapidly into high-impact domains:

- finance,
- healthcare,
- pricing,
- compliance,
- and policy enforcement.

In these contexts, **confidence without trust is dangerous**.

Regulatory pressure, reputational risk, and ethical scrutiny are converging on one requirement: **defensible data use**.

Organizations that manage trust explicitly gain:

- higher automation ceilings,
- safer AI deployment,
- and strategic advantage.

Those that do not will be forced to slow down or roll back.

The key takeaway

AI does not fail because it lacks intelligence.

It fails because it lacks **trust boundaries**.

Data trust is not a feature, a policy, or a compliance checkbox.
It is the **structural prerequisite** for AI systems that can operate responsibly, defensibly, and at scale.

All rights reserved.

Table of contents

Intended audience 1
Purpose of this document 2
Why Data Trust Is the Defining AI Constraint 4
Defining Data Reliability 8
Context as a Reliability Multiplier 14
Data Movement and Reliability Degradation 18
Human Intervention and Control Points 21
Corrections, Overrides, and Trust Recovery 24
Decision Boundaries by Reliability Level 27
Organizational Readiness 31
Assessing Your Organization's Trust Readiness 35
Appendix A — Definitions & Glossary 38
About the Author 41

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Intended audience

This document is intended for professionals and organizations that **create, manage, consume, or act upon data in decision-critical, data-driven, contexts**.

It is specifically written for:

- **Executive leadership (C-level, board members)**
  Responsible for strategic decisions, governance, and risk management where data trust directly impacts organizational outcomes.
- **Data leaders and architects**
  Including heads of data, data platform owners, enterprise architects, and analytics leads who design and maintain data flows across systems and domains.
- **Compliance, risk, and governance professionals**
  Involved in regulatory compliance, auditability, internal controls, and accountability frameworks (e.g. financial reporting, AI governance, or regulated decision-making).
- **AI and advanced analytics practitioners**
  Who rely on data inputs for model training, inference, and automation, and must understand the limits of what data can safely be used for.
- **Domain experts and data owners**
  Who provide, curate, or validate data within specific business contexts and are accountable for its correct interpretation.

What this document is *not*

This document is **not**:

- a product manual,
- a technical implementation guide,
- a vendor-specific solution description,
- or a marketing or sales narrative.

Instead, it is a **conceptual and operational framework** designed to help organizations:

- reason clearly about data trust,
- define reliability levels consistently,
- and assess their own data landscape with rigor.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Assumed background

The document assumes that the reader:

- has basic familiarity with data systems and organizational processes,
- understands that data moves, transforms, and accumulates context over time,
- and is involved—directly or indirectly—in decisions where data quality alone is insufficient.

No prior expertise in data science, statistics, or machine learning is required.
However, the reader is expected to engage with **governance, accountability, and risk trade-offs** at an organizational level.

Goal for the reader

After reading this document, the reader should be able to:

- distinguish clearly between reliability, quality, accuracy, and confidence,
- identify trust boundaries within their own data flows,
- recognize where context is gained or lost,
- and determine where human intervention, controls, or policy are required.

The ultimate goal is **not** to prescribe specific tools or architectures, but to establish a shared language and decision framework for trustworthy data use.

Purpose of this document

The purpose of this document is to establish a **clear, shared framework** for understanding, assessing, and managing data trust in modern organizations.

As data increasingly drives decisions, automation, and AI systems, organizations face a fundamental challenge:
they rely on data whose **trustworthiness is often assumed, but rarely made explicit**.

This document exists to address that gap.

What this document aims to achieve

This document is designed to help organizations:

- Develop **awareness** of data trust as a distinct and critical concern
- Establish **clear definitions** for reliability, context, and trust boundaries
- Understand how data movement, transformation, and enrichment affect trust
- Identify where **human judgment** is required in otherwise automated systems
- Define **decision boundaries** based on data reliability
- Assess organizational readiness to operate with trusted data

The goal is not to prescribe specific tools, architectures, or vendors, but to provide a **conceptual and operational foundation** that can be applied across technologies, sectors, and organizational models.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Why this document is necessary

Many organizations already invest heavily in:

- data platforms,
- analytics,
- AI models,
- and governance initiatives.

Yet they still struggle to answer basic trust questions, such as:

- *Why do we trust this data?*
- *Who is accountable if this decision is challenged?*
- *Under what conditions does this data stop being valid?*

Without explicit answers, trust decisions become:

- implicit,
- inconsistent,
- and dependent on individuals rather than systems.

This document makes those decisions **explicit, discussable, and governable**.

What this document does *not* attempt to do

This document does **not**:

- define universal trust scores,
- provide legal advice,
- replace regulatory frameworks,
- or claim to eliminate risk.

Trust is contextual, domain-specific, and ultimately a matter of responsibility.
No framework can remove the need for judgment.

Instead, this document aims to:

- surface assumptions,
- clarify trade-offs,
- and enable informed decision-making.

Intended outcome

After engaging with this document, an organization should be able to:

- articulate its own trust assumptions,
- recognize where those assumptions break down,
- identify gaps between desired automation and actual trust readiness,
- and take deliberate steps toward more defensible data use.

The ultimate outcome is **not certainty**, but **control**.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

A shared language for trust

Perhaps most importantly, this document provides a **shared vocabulary**.

Without shared language:

- trust concerns remain vague,
- disagreements become personal,
- and risk is addressed reactively rather than structurally.

By establishing common concepts and boundaries, this document enables:

- productive internal dialogue,
- alignment across functions,
- and more responsible use of data and AI.

Final note

This document is intentionally practical, normative, and experience-driven.
It reflects patterns observed across industries and decision contexts where data trust is not optional, but foundational.

Its purpose is simple:

To help organizations use data in ways they can stand behind—today, tomorrow, and under scrutiny.

Why Data Trust Is the Defining AI Constraint

Artificial intelligence is often discussed in terms of models, compute, and data volume. Public discourse focuses on architectural breakthroughs, scaling laws, and performance benchmarks. While these elements matter, they are **not** what ultimately determines whether AI systems can be deployed responsibly, sustainably, and at scale.

The defining constraint for AI is **data trust**.

AI does not reason about truth — it operationalizes input

AI systems do not understand truth, intent, or accountability. They transform input data into outputs based on statistical patterns, learned representations, or predefined rules.

As a result:

- AI does not question whether data *should* be used.
- It does not understand where data comes from.
- It cannot assess whether acting on data is justified.

AI systems therefore **amplify the properties of their inputs**.
They do not correct unreliable data — they scale it.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Confidence is not trust

Modern AI systems often produce outputs accompanied by confidence scores or probabilities. This creates a dangerous illusion: that uncertainty has been managed simply because it is quantified.

In reality:

- confidence expresses internal model certainty,
- trust expresses external defensibility.

A system can be highly confident while being fundamentally untrustworthy, because:

- the source is unknown,
- the context is missing,
- or accountability is absent.

Treating confidence as a proxy for trust is one of the most common and consequential errors in AI deployment.

The automation paradox

AI promises efficiency, consistency, and scale. These benefits are most attractive in domains where:

- decisions are frequent,
- outcomes are impactful,
- and human intervention is costly.

Ironically, these are also the domains where **trust requirements are highest**:

- finance and credit,
- healthcare,
- pricing and allocation,
- compliance and enforcement,
- risk and fraud detection.

The paradox is this:

The more a decision is automated, the more critical it becomes to justify the data behind it.

Without explicit trust constraints, automation increases both **speed and blast radius** of failure.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Why better models do not solve the problem

There is a persistent belief that sufficiently advanced models can compensate for weak data foundations.

They cannot.

No model can:

- reconstruct missing provenance,
- infer contractual or legal context,
- assign responsibility retroactively,
- or make data defensible under challenge.

Model sophistication increases **confidence**, not **legitimacy**.

As models become more capable, the cost of trust failure rises rather than falls.

Traditional data governance is not enough

Most organizations already have data governance practices:

- access controls,
- data quality metrics,
- lineage diagrams,
- validation rules.

These practices answer operational questions:

- Is the data complete?
- Is it consistent?
- Who can access it?

They do **not** answer trust questions:

- Should this data be used for this decision?
- What assumptions does this data carry?
- Who is accountable if the outcome is challenged?
- Where does responsibility shift as data moves?

AI systems require answers to these trust questions. Without them, governance remains superficial.

Trust is a gating mechanism, not a property

Data trust should not be treated as a score to be optimized.
It should be treated as a **gate** that controls what actions are allowed.

Trust-aware systems ask:

- What decisions are permitted at this reliability level?
- When must human judgment intervene?
- When must automation stop?

This framing turns trust from an abstract principle into an **operational constraint**.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Regulatory pressure makes trust unavoidable

Emerging regulation (e.g. AI governance, financial oversight, sector-specific compliance) increasingly demands:

- explainability,
- auditability,
- accountability,
- and defensibility.

These requirements cannot be met retroactively.

Organizations that do not manage data trust proactively will be forced to:

- limit AI usage,
- roll back automation,
- or accept unacceptable legal and reputational risk.

Trust is no longer optional; it is becoming **structural**.

Trust determines the ceiling of AI impact

Organizations often ask:

“How far can we go with AI?”

The real answer is:

As far as your data trust allows.

Without explicit trust management:

- AI remains confined to low-risk, advisory use cases.
- High-impact decisions remain manual.
- Scaling stalls under scrutiny.

Organizations that manage trust deliberately unlock:

- deeper automation,
- higher-value use cases,
- and sustained competitive advantage.

The real bottleneck

The limiting factor for AI is not:

- model accuracy,
- inference latency,
- or infrastructure cost.

It is the inability to answer, consistently and rigorously:

Which data may be used for which decisions — and why?

Until organizations can answer this question, AI adoption will remain constrained by risk rather than ambition.

© 2026 Martijn Wiggers
This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Key takeaway

AI does not fail because it lacks intelligence.

It fails because it lacks **trust boundaries**.

Data trust is therefore not an AI feature, a compliance add-on, or an ethical afterthought.
It is the **precondition** for AI systems that can operate responsibly, defensibly, and at scale.

Defining Data Reliability

Reliability vs Quality vs Accuracy vs Confidence

In discussions about data and AI, terms such as *reliability*, *quality*, *accuracy*, and *confidence* are often used interchangeably. While related, these concepts describe **fundamentally different properties** of data. Failing to distinguish between them leads to incorrect assumptions, unsafe automation, and ultimately decisions that cannot be defended.

This section defines these concepts precisely and explains how they relate to one another.

Reliability — Can this data be trusted *at all*?

Reliability describes the extent to which data is **defensible, traceable, and accountable** within a decision-making context.

Data is reliable when:

- its origin is known,
- the method of acquisition is explainable,
- and it is clear **who is responsible** for its correctness.

Crucially, reliability is **not** primarily about whether data is factually correct, but about:

whether one is justified in acting as if the data were correct.

Data may be factually accurate yet unreliable, for example when:

- the source is unknown,
- critical context is missing,
- or no party can be held accountable for its validity.

Reliability is therefore a **prerequisite** for any form of automated, regulatory, or high-impact decision-making.

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Accuracy — Is the value correct?

Accuracy describes how close a data value is to the *actual real-world value*.

Examples include:

- a temperature reading of 20.1°C when the true value is 20.0°C,
- a reported transaction price of €500,000 when the actual amount was €510,000.

Accuracy is:

- **quantitative**
- **measurable**
- often expressed through error margins or deviations

However, accuracy says **nothing** about:

- data provenance,
- contextual relevance,
- legitimacy,
- or responsible use.

Highly accurate data can still be unusable if its origin or use cannot be justified.

Quality — Does the data meet its intended purpose?

Data quality describes the degree to which data is fit for a **specific intended use**.

Quality is therefore **context-dependent** and may include dimensions such as:

- completeness
- consistency
- timeliness
- conformity to standards
- structural correctness

A dataset may be high-quality for reporting purposes but low-quality for real-time decision-making. Quality is therefore always **relative to a goal or application**.

Importantly, quality is often measured operationally, while reliability is fundamentally a **governance concern**.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Confidence — How certain do we *feel* about the data?

Confidence represents a **degree of certainty or probability** assigned to a data point or a conclusion derived from data.

In modern AI systems, confidence is often:

- statistically calculated,
- derived from model outputs,
- expressed as a probability score or likelihood

Confidence may exist:

- without full context,
- without verifiable provenance,
- without explicit accountability

As such, confidence is **not a property of the data itself**, but of:

- a model,
- an estimation,
- or an interpretation.

It is entirely possible to have:

- high confidence in unreliable data,
- low confidence in highly reliable data.

Confidence must therefore **never be treated as a substitute for reliability**.

Why these distinctions matter

A fundamental issue in many data- and AI-driven organizations is that:

- accuracy is used as a proxy for reliability,
- quality is confused with trustworthiness,
- confidence is interpreted as legitimacy.

This leads to situations where:

- decisions are made based on data that cannot be defended,
- automation is deployed without a clear accountability chain,
- risks only become visible after damage has occurred.

Clearly separating these concepts makes it possible to:

- constrain decision-making appropriately,
- assign responsibility explicitly,
- deploy automation safely and at scale.

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Summary

Concept	Primary question it answers
Reliability	*May we rely on this data for decisions?*
Accuracy	*Is the value correct?*
Quality	*Is the data fit for this specific purpose?*
Confidence	*How certain is a model or human about the outcome?*

Reliability is **foundational**.
Accuracy, quality, and confidence are **supporting dimensions**.

Without reliability, none of the others are sufficient.

A Practical Reliability Scale for data

To reason meaningfully about data trust, organizations need more than a binary distinction between "trusted" and "untrusted" data. In practice, data exists on a **spectrum of reliability**, shaped by provenance, verifiability, contractual guarantees, and contextual completeness.

This section introduces a **practical reliability scale** that can be applied consistently across industries, domains, and regulatory environments. The scale is intentionally grounded in a simple but demanding question:

Can this data be defended under scrutiny, and if so—by whom, and with what context?

Reliability is about defensibility, not belief

Reliability is often conflated with confidence, accuracy, or popularity. None of these are sufficient.

Reliable data is data whose **claims can be substantiated**:

- with evidence,
- by identifiable parties,
- under defined conditions,
- and within a known context.

The ultimate stress test for reliability is not internal agreement or system confidence, but **external challenge**—for example in audits, disputes, or court proceedings. Not because all data will end up in court, but because this forces clarity about evidence, ownership, and accountability.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

The Reliability Scale

The following six-level scale describes increasing levels of data reliability, from unverifiable signals to fully defensible facts with context.

Level 1 — Unknown or Unverifiable Origin

Reliability: Minimal

- The origin of the data is unknown or undocumented.
- No party can credibly attest to its correctness.
- No evidence, audit trail, or contractual responsibility exists.

Examples:

- Scraped data without provenance
- Anonymous reports or undocumented exports
- Legacy datasets with missing metadata

This data may be useful for exploration or hypothesis generation, but **must not** be used for decisions with material impact.

Level 2 — Unprovable Assertions

Reliability: Low

- A known source exists, but the data cannot be proven or defended.
- Assertions are based on statements, beliefs, or informal communication.
- No enforceable obligation to provide evidence.

Examples:

- Rumours, estimates, or hearsay
- Internal assumptions without documentation
- “We believe this to be correct” data

This level supports narrative or directional insight, but not operational or financial decisions.

Level 3 — Contractually Asserted Truth

Reliability: Medium

- A third party asserts correctness under a contract or SLA.
- The asserting party claims the ability to prove correctness if required.
- Context may be incomplete or implicit.

Examples:

- Vendor-provided datasets under SLA
- Partner system integrations
- Certified reports without raw evidence attached

Reliability depends on **legal enforceability**, not direct evidence possession. Trust is delegated.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Level 4 — Independently Verifiable Facts

Reliability: High

- The data can be proven with direct evidence.
- Supporting documents, transactions, or records are accessible.
- Proof does not depend on trust in a third party alone.

Examples:

- Transaction logs with matching confirmations
- Signed contracts and recorded events
- Traceable system-of-record outputs

At this level, data can withstand audit and dispute—but may still lack full situational meaning.

Level 5 — Verifiable Facts with Explicit Context

Reliability: Very High

- All criteria of Level 4 are met.
- Relevant **context** is explicitly captured and preserved.
- Interpretative risks are minimized.

Context includes:

- Conditions under which data was produced
- Exceptional circumstances (e.g. forced sale, emergency override)
- Temporal, legal, or procedural qualifiers

This level supports **high-impact decisions**, valuations, automation, and regulatory reporting.

Level 6 — Authoritative or Mandated Truth

Reliability: Maximum (Context-dependent)

- Data is issued or validated by a recognized authority.
- Legal or regulatory frameworks mandate acceptance.
- In some domains, this exceeds self-held evidence.

Examples:

- Government registries
- Court rulings
- Official land, identity, or corporate records

Notably, authority does not eliminate the need for context. Even mandated truth can be misinterpreted if stripped of qualifying information.

© 2026 Martijn Wiggers
This document provides conceptual guidance.
It does not constitute legal or regulatory advice.

All rights reserved.

Reliability is multi-dimensional

Several clarifications are critical:

- **Reliability can differ by level**
  Source-level reliability may differ from record- or field-level reliability.
- **Context is not optional**
  Data without context may be provable yet misleading.
- **Self vs third party matters**
  Who can prove the data—and under what obligation—materially affects reliability.
- **Higher reliability does not mean higher accuracy**
  It means higher defensibility.

Why this scale matters

Without a shared reliability language:

- organizations over-automate on weak data,
- underutilize strong data,
- and fail to recognize where human judgment is still required.

This scale enables:

- explicit decision boundaries,
- clearer governance,
- and honest conversations about where trust truly ends.

In the next section, we examine how **context acts as a reliability multiplier**—and why ignoring it leads to systematically wrong conclusions, even when the underlying facts are correct.

Context as a Reliability Multiplier

Facts do not exist in isolation.
Without context, even correct data can be misleading, unsafe, or outright dangerous when used for decision-making.

Context does not replace reliability — it **multiplies or constrains** it.

What is context?

In the context of data trust, **context** refers to information that explains *how, why, when, and under which conditions* data came into existence.

Context may include:

- legal or contractual conditions,
- temporal constraints,
- exceptional circumstances,
- domain-specific meaning,
- intent behind a transaction or record,
- known limitations or exclusions.

© 2026 Martijn Wiggers
This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Context answers the question:

Under which assumptions is this data valid?

Facts without context are incomplete truths

A data value may be factually correct while still being unusable or misleading.

For example:

- a transaction price is correct, but reflects a forced sale,
- a measurement is accurate, but taken during an exceptional event,
- a contract exists, but under non-standard terms,
- a dataset is complete, but only for a restricted population.

Without context, such data invites **overgeneralization**.

This is not a data quality issue.
It is a **reliability issue**.

Context multiplies reliability — it does not add to it

Context should not be treated as an additional attribute layered on top of reliability. Instead, it **modifies how reliability should be interpreted**.

Conceptually:

- High reliability **without context** → limited, conditional trust
- High reliability **with context** → defensible, actionable trust
- Low reliability **with context** → bounded, cautious use
- Low reliability **without context** → non-actionable

Context therefore acts as a **multiplier**, not a score.

It determines:

- which conclusions may be drawn,
- which decisions are allowed,
- and where boundaries must be enforced.

Context can increase or constrain trust

Context does not always increase trust.
In many cases, it **constrains** it.

For example:

- Knowing that a price resulted from an execution sale lowers its suitability as a market reference.
- Knowing that data was collected during a system outage limits its representativeness.
- Knowing that a dataset excludes certain populations restricts its applicability.

Context does not weaken data — it **prevents misuse**.

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Context exists at multiple levels

Context can apply at different granularities:

- **Source-level context**
  Applies to all data from a source (e.g. regulatory scope, authority, systemic bias).
- **Record-level context**
  Applies to specific entries (e.g. exceptional transactions, overrides, special conditions).
- **Field-level context**
  Applies to individual attributes (e.g. estimated values, derived fields, provisional status).

Failing to distinguish these levels leads to either:

- over-restricting useful data,
- or over-trusting critical exceptions.

Context is fragile and easily lost

Context is often:

- implicit,
- informal,
- or held only in human understanding.

As data moves through systems, context is frequently:

- stripped away,
- collapsed into comments,
- or replaced by assumptions.

Once context is lost, it **cannot reliably be reconstructed**.

This loss is a primary driver of trust decay.

Context must be preserved explicitly

Trustworthy systems treat context as:

- first-class information,
- preserved alongside data,
- and visible to decision-makers.

This includes:

- documenting assumptions,
- marking exceptional conditions,
- preserving temporal relevance,
- and preventing context collapse during aggregation or transformation.

Context that exists only “in people’s heads” does not scale.

© 2026 Martijn Wiggers
This document provides conceptual guidance.
It does not constitute legal or regulatory advice.

All rights reserved.

Context and accountability

Context determines responsibility.

When context is missing:

- accountability becomes ambiguous,
- blame shifts from decisions to systems,
- and defensibility collapses.

Explicit context makes it possible to answer:

- *Who knew what, and when?*
- *Which assumptions were accepted?*
- *Why was this data considered valid at the time?*

These are trust questions, not technical ones.

Context-aware decision boundaries

Decision boundaries must consider both:

- reliability level,
- and contextual constraints.

A dataset may be reliable enough for:

- internal analysis,

but not for:

- pricing,
- enforcement,
- or automated action.

Context determines **where reliability applies and where it does not**.

Key takeaway

Reliability determines *whether* data may be trusted.
Context determines *how far that trust extends*.

Without context:

- facts mislead,
- automation overreaches,
- and defensibility disappears.

Context is therefore not optional metadata.
It is the **multiplier that makes reliable data safe to use**.

© 2026 Martijn Wiggers
This document provides conceptual guidance.
It does not constitute legal or regulatory advice.

All rights reserved.

Data Movement and Reliability Degradation

Data Movement and Trust Decay

Data does not remain static. It moves between systems, is copied, transformed, enriched, aggregated, filtered, inferred, and merged. Each of these movements affects the **reliability** of the data, often in subtle but consequential ways.

This gradual loss or alteration of trustworthiness is referred to here as **trust decay**.

Data movement is not neutral

Every data movement introduces at least one of the following risks:

- loss of provenance
- loss of contextual information
- dilution of accountability
- increased ambiguity about correctness
- separation between data and its original purpose

Even when data values remain unchanged, the **meaning and defensibility** of those values may degrade.

For example:

- copying data to another system may remove information about who verified it,
- exporting data to a report may detach it from its original contractual or legal context,
- aggregating data may obscure critical outliers or exceptional conditions,
- transforming data may implicitly introduce assumptions that are no longer visible.

Trust decay is therefore not necessarily caused by errors, but by **distance**—distance from source, from context, and from responsibility.

Common trust-decaying operations

The following operations typically reduce reliability unless explicitly controlled:

Copying and replication

Data duplicated across systems often loses:

- ownership clarity,
- update guarantees,
- and explicit responsibility.

Without strong controls, copies quickly diverge from their authoritative source.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Transformation and normalization

Transformations such as unit conversion, normalization, categorization, or formatting may:

- embed assumptions,
- hide original values,
- or make reversibility impossible.

While often necessary, transformations require explicit documentation to preserve trust.

Aggregation and summarization

Aggregation reduces detail and removes variance:

- averages hide distributions,
- totals hide individual transactions,
- summaries hide exceptions.

As a result, aggregated data typically has **lower reliability for decision-making** than the underlying records, even if statistically correct.

Enrichment and joins

Enriching data with external sources increases informational value but introduces:

- dependency on third-party reliability,
- mixed accountability,
- and blended trust levels.

The resulting dataset is only as reliable as its **weakest critical component**.

Derivation and inference

Derived values—such as scores, classifications, or predictions—are **not facts**, even when computed deterministically.

They represent:

- assumptions,
- models,
- or interpretations.

Derived data inherits reliability from:

- its inputs,
- the transparency of the method,
- and the accountability for the derivation logic.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Trust decay is cumulative

Trust decay compounds over time and across steps.

A dataset that passes through multiple systems, transformations, and owners may become:

- operationally useful,
- statistically valid,
- and technically correct,

while simultaneously becoming **non-defensible** for high-impact decisions.

Importantly, trust decay is often **invisible** in technical systems.
Pipelines may run flawlessly while trust silently erodes.

Preventing uncontrolled trust decay

Trust decay cannot be eliminated, but it **can be managed**.

Effective strategies include:

- explicit trust levels attached to data at appropriate granularity,
- preservation of provenance and context metadata,
- clear accountability at each transformation step,
- constraints on how data of certain reliability levels may be used,
- and human validation where automated guarantees are insufficient.

Crucially, organizations must accept that:

Not all data movements are equal, and not all outputs are fit for all decisions.

Trust decay vs data quality degradation

Trust decay should not be confused with data quality issues.

- Data quality degradation refers to errors, missing values, or inconsistencies.
- Trust decay refers to **loss of defensibility**, even when data remains correct.

High-quality data can still be untrustworthy if its lineage, context, or responsibility chain is unclear.

Key takeaway

Data reliability is **not preserved by default**.

Every movement, transformation, and enrichment step must be treated as a **trust decision**, not merely a technical operation.

Organizations that fail to manage trust decay inevitably end up with data that:

- looks authoritative,
- powers dashboards and models,
- but cannot safely support the decisions it is used for.

All rights reserved.

Human Intervention and Control Points

Human intervention is often discussed as a fallback mechanism for automated or AI-driven systems. In practice, this framing is incorrect and dangerous.

Human involvement is not a last resort.
It is a **structural control mechanism** that must be deliberately designed into data and decision pipelines.

Why human intervention is unavoidable

No system—regardless of its technical sophistication—can fully account for:

- missing context,
- exceptional circumstances,
- normative judgment,
- or evolving interpretations of correctness.

Certain aspects of trust cannot be automated because they involve:

- intent,
- responsibility,
- and accountability.

As a result, any system that claims to eliminate the need for human judgment is implicitly denying responsibility.

Control points vs ad-hoc intervention

Human intervention must occur at **defined control points**, not as improvised exception handling.

Ad-hoc intervention:

- introduces inconsistency,
- hides accountability,
- and creates undocumented trust changes.

Well-defined control points:

- are intentional,
- repeatable,
- and auditable.

They define *where* and *why* humans may intervene—not merely *that* they can.

© 2026 Martijn Wiggers
This document provides conceptual guidance.
It does not constitute legal or regulatory advice.

All rights reserved.

Typical human control points in data flows

The following control points are commonly required in trustworthy systems.

1. Source onboarding and classification

When a new data source is introduced, humans must:

- assess provenance and legitimacy,
- evaluate contractual or legal guarantees,
- assign an initial reliability level,
- and document assumptions.

This step cannot be automated because it involves interpretation of external agreements, authority, and intent.

2. Context validation

Context often exists outside structured data:

- legal conditions,
- exceptional events,
- business-specific meaning,
- or domain knowledge.

Human validation is required to:

- confirm whether context still applies,
- identify changes that invalidate assumptions,
- or flag data as conditional or exceptional.

3. Trust escalation approval

Escalating data usage—for example from advisory to automated decisions—must require explicit human approval.

This ensures that:

- reliability thresholds are consciously accepted,
- risk trade-offs are understood,
- and responsibility is clearly assigned.

Automation without explicit escalation approval is a governance failure.

© 2026 Martijn Wiggers
This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

4. Exception handling and overrides

When systems encounter anomalies or contradictions, humans may:

- override automated outcomes,
- correct data,
- or suspend processing.

Crucially, overrides must:

- be intentional,
- be logged with justification,
- and result in explicit trust state changes.

Silent overrides undermine trust rather than restore it.

5. Corrections and trust recovery

When data is corrected, humans must:

- determine the scope of the correction,
- assess downstream impact,
- and decide whether trust can be restored.

Not all corrections fully restore reliability.
Some permanently lower defensibility due to lost evidence or ambiguity.

Human intervention is not a quality fix

Human involvement cannot:

- retroactively create provenance,
- eliminate missing historical context,
- or convert unverifiable data into facts.

Human control exists to **govern decisions**, not to cosmetically improve data.

Treating human review as a substitute for trust is a common but flawed approach.

Avoiding human bottlenecks

While human intervention is necessary, it must be:

- limited to defined points,
- proportional to decision impact,
- and supported by clear criteria.

Effective systems:

- automate low-risk decisions,
- require human input only where trust boundaries are crossed,
- and provide humans with sufficient context to act meaningfully.

Overusing human intervention leads to bottlenecks; underusing it leads to systemic risk.

All rights reserved.

Accountability remains human

Regardless of automation level:

Accountability for decisions always rests with humans.

Systems may recommend, optimize, or execute actions—but responsibility cannot be delegated to software.

Human control points exist to make this accountability explicit and enforceable.

Key takeaway

Human intervention is not a weakness of data-driven or AI systems.

It is a **necessary design principle** for any system that aims to operate:

- responsibly,
- defensibly,
- and at scale.

Organizations that fail to design explicit human control points do not remove risk—they merely hide it.

Corrections, Overrides, and Trust Recovery

No data system is static. Errors occur, assumptions change, context evolves, and new information becomes available. In trustworthy systems, the question is not **whether** corrections and overrides happen, but **how** they are handled.

This section defines how corrections affect reliability and under what conditions trust can—or cannot—be restored.

Corrections do not automatically restore trust

A common misconception is that correcting a value restores its original trust level.

This is not always true.

A correction may:

- fix an incorrect value,
- improve accuracy,
- or resolve an inconsistency,

while simultaneously **reducing reliability**, because:

- the original state can no longer be proven,
- historical context is altered,
- or intent becomes ambiguous.

Reliability depends not only on correctness, but on **defensibility over time**.

© 2026 Martijn Wiggers
This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Types of corrections and their impact

Not all corrections affect trust equally.

1. Evidenced corrections

Evidenced corrections are supported by:

- original source documents,
- verifiable transactions,
- or authoritative records.

Characteristics:

- provenance is preserved,
- justification is explicit,
- responsibility is clear.

These corrections may **partially or fully restore trust**, depending on context.

2. Interpretive corrections

Interpretive corrections occur when:

- new understanding emerges,
- classification rules change,
- or context is reinterpreted.

Characteristics:

- correctness is subjective,
- multiple interpretations may be valid,
- defensibility depends on documentation.

These corrections usually **lower reliability**, even if accuracy improves.

3. Speculative or inferred corrections

Speculative corrections are made when:

- original evidence is missing,
- data is reconstructed,
- or assumptions replace facts.

Characteristics:

- provenance is weak or absent,
- accountability is diluted,
- reversibility is limited.

Such corrections **cannot restore trust** and often permanently lower it.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Overrides as trust-altering actions

Overrides are intentional deviations from automated or predefined outcomes.

They are sometimes necessary—but always consequential.

Overrides:

- must be explicit,
- must be justified,
- and must be logged.

An override changes the trust state because:

- it introduces human judgment,
- bypasses automated guarantees,
- and alters accountability dynamics.

Undocumented overrides are indistinguishable from errors.

Trust recovery is conditional, not guaranteed

Trust recovery depends on three factors:

1. **Evidence availability**
   Can the original claim still be proven?
2. **Context preservation**
   Is the original meaning and intent still intact?
3. **Accountability continuity**
   Is responsibility clearly assigned before and after correction?

If any of these are missing, trust recovery is limited.

In some cases, the correct outcome is to **accept permanent trust degradation**.

Temporal integrity matters

Corrections affect not only current data, but historical decisions.

Organizations must be able to answer:

- What decisions were made based on the original data?
- Were those decisions justified *at the time*?
- Does the correction invalidate past actions?

Trustworthy systems preserve:

- historical states,
- correction timestamps,
- and decision context.

Rewriting history erodes trust more than acknowledging error.

© 2026 Martijn Wiggers
This document provides conceptual guidance.
It does not constitute legal or regulatory advice.

All rights reserved.

Avoiding trust inflation

One of the most dangerous patterns is **trust inflation**:

- silently restoring high trust levels after correction,
- treating corrected data as equivalent to original facts,
- or masking uncertainty to maintain automation.

Trust inflation creates false confidence and systemic risk.

Conservative trust recovery is always preferable to optimistic assumptions.

Governance implications

Organizations must define:

- who is allowed to correct data,
- under what conditions overrides are permitted,
- how trust levels change as a result,
- and when escalation is required.

Without explicit rules, corrections become discretionary—and trust becomes arbitrary.

Key takeaway

Corrections fix values.
Overrides fix outcomes.
Neither automatically fixes trust.

Trust recovery is a **governed process**, not a technical side effect.

Organizations that treat corrections as purely operational events will inevitably undermine the defensibility of their data and decisions.

Decision Boundaries by Reliability Level

Not all data may be used for all decisions.
Treating data as universally actionable—regardless of its reliability—creates systemic risk.

Decision boundaries define **which types of actions are permissible** at each level of data reliability.
These boundaries ensure that organizations:

- do not over-automate on insufficiently reliable data,
- preserve accountability,
- and maintain defensibility as decisions scale.

© 2026 Martijn Wiggers
This document provides conceptual guidance.
It does not constitute legal or regulatory advice.

All rights reserved.

Why decision boundaries are necessary

In human organizations, trust boundaries exist implicitly:

- senior decisions require “verified numbers,”
- early-stage insights are labeled “directional,”
- and legally binding actions require documented proof.

AI systems, however, lack intuition.
Without explicit boundaries, they will treat all inputs as equally actionable.

Decision boundaries make trust **explicit and enforceable**.

Reliability as a gating mechanism

Reliability should function as a **gate**, not a score.

Rather than asking:

“How good is this data?”

The relevant question is:

“What are we allowed to do with this data?”

This reframes reliability from a descriptive metric into an **operational control**.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Example decision boundaries by reliability level

The table below illustrates how reliability levels map to permissible decision types.
(Exact thresholds and definitions will vary by organization and domain.)

Reliability Level	Description (Illustrative)	Allowed Decisions	Prohibited Decisions
Level 1	Rumor, unverified, unknown source	Exploratory analysis, hypothesis generation, internal brainstorming	Automated decisions, external reporting, financial or legal actions
Level 2	Indicative, low confidence, partial context	Trend spotting, early warning signals, non-binding recommendations	Customer-facing actions, enforcement, irreversible decisions
Level 3	Contractually asserted by third party	Advisory decisions, human-in-the-loop workflows, internal optimization	Fully automated decisions, regulatory filings
Level 4	Verifiable, accountable, with evidence	Operational decisions, customer-impacting actions with oversight	High-stakes automation without review
Level 5	Fully defensible, provable with context	Automated decision-making, regulatory reporting, financial settlement, adding to showcase	—

Human-in-the-loop as a boundary, not a fix

Human review is often introduced as a blanket safety mechanism.
In reality, it only works **within defined trust boundaries**.

Human intervention can:

- validate assumptions,
- add missing context,
- approve exceptions.

It cannot:

- retroactively make unreliable data reliable,
- reconstruct missing provenance,
- or eliminate accountability gaps.

Human-in-the-loop must therefore be treated as a **controlled boundary**, not a universal remedy.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Escalation and degradation

Decision boundaries should support:

- **escalation** when higher-reliability data becomes available,
- **degradation** when reliability drops due to trust decay.

For example:

- a recommendation may escalate to automation once reliability thresholds are met,
- an automated decision may degrade to advisory mode if data quality or provenance is compromised.

This dynamic behavior is essential in real-world systems.

Why boundaries must be enforced, not documented

Decision boundaries that exist only in policy documents are ineffective.

To be meaningful, boundaries must be:

- enforced in systems,
- visible to decision-makers,
- and auditable after the fact.

Without enforcement:

- trust classifications become decorative,
- risk silently accumulates,
- and organizations lose control over AI-driven outcomes.

Key takeaway

Reliability is not about measuring trust—it is about **controlling action**.

Decision boundaries translate abstract trust concepts into:

- concrete permissions,
- enforceable constraints,
- and defensible behavior.

Without decision boundaries, AI systems will always operate beyond what data trust can justify.

© 2026 Martijn Wiggers
This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Organizational Readiness

Data trust is not achieved through technology alone.
It requires an organization to be **structurally prepared** to take responsibility for how data is created, transformed, interpreted, and acted upon.

Many organizations believe they are “data-driven,” yet lack the organizational foundations required to use data **reliably** and **defensibly**—especially in automated or AI-driven contexts.

This section outlines what organizational readiness for data trust actually entails.

Trust readiness is not maturity

Organizational readiness should not be confused with data maturity models.

An organization may have:

- modern data platforms,
- advanced analytics,
- and sophisticated AI capabilities,

while still being **unready** to operate with trusted data.

Readiness is not about how advanced systems are.
It is about whether the organization can **answer hard questions** when decisions are challenged.

Core organizational prerequisites

An organization is ready to work with data trust when the following conditions are met.

1. Clear ownership and accountability

Every critical dataset must have:

- a clearly identified owner,
- explicit responsibility for correctness and context,
- and authority to approve or restrict usage.

Without ownership:

- reliability cannot be asserted,
- corrections cannot be enforced,
- and accountability collapses under scrutiny.

Ownership is an organizational role, not a technical attribute.

© 2026 Martijn Wiggers
This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

2. Explicit decision responsibility

Organizations must be able to answer:

Who decided that this data may be used for this decision?

This requires:

- defined decision-makers,
- documented approval thresholds,
- and clarity on who bears responsibility if outcomes are contested.

In trustworthy systems, **decisions are traceable to people**, not just pipelines.

3. Willingness to constrain automation

Organizational readiness includes the ability to say:

“We will not automate this yet.”

This requires:

- acceptance of limits,
- resistance to premature optimization,
- and recognition that some decisions require human judgment.

Organizations that treat automation as an unquestioned goal will inevitably exceed their trust boundaries.

4. Tolerance for uncertainty and degradation

Trusted organizations acknowledge that:

- data reliability can decrease,
- context can be lost,
- and systems can degrade.

Readiness includes:

- fallback procedures,
- escalation paths,
- and acceptance of degraded modes of operation.

Organizations that assume trust is permanent will be surprised when it fails.

© 2026 Martijn Wiggers
This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

5. Cross-functional alignment

Data trust cannot be owned by a single department.

It requires alignment between:

- business leadership,
- data and engineering teams,
- legal and compliance,
- risk management,
- and domain experts.

When these functions operate in isolation, trust decisions become fragmented and inconsistent.

Cultural indicators of readiness

Certain cultural signals strongly correlate with trust readiness:

- Decisions are routinely challenged constructively.
- Assumptions are documented, not implied.
- “Directionally correct” is distinguished from “defensible.”
- Exceptions are tracked, not normalized.
- Discomfort with uncertainty is tolerated.

Organizations lacking these traits often rely on confidence and momentum rather than trust.

Common readiness anti-patterns

Organizations are **not ready** when they exhibit patterns such as:

- “The dashboard says so” as a justification.
- Treating data quality metrics as proof of trust.
- Relying on models to compensate for weak data foundations.
- Diffuse responsibility (“the system decided”).
- Governance documents without enforcement mechanisms.

These patterns indicate that trust is assumed, not managed.

Readiness is incremental

Organizational readiness is not binary.

Most organizations will:

- start with limited trust-aware decisions,
- gradually formalize ownership and boundaries,
- and expand scope as confidence in governance grows.

Attempting full-scale trust enforcement without organizational readiness may result in resistance or superficial compliance.

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Key takeaway

Data trust is an **organizational capability**, not a technical feature.

Technology can support trust, but only organizations that:

- accept accountability,
- enforce boundaries,
- and align decision-making,

are truly ready to operate with reliable data—especially in AI-driven environments.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Assessing Your Organization's Trust Readiness

A Practical Self-Assessment

Before implementing trust-aware systems, organizations must assess whether they are **organizationally prepared** to do so. Trust readiness is not a technical milestone; it is a structural and cultural condition.

This self-assessment is designed to help organizations identify gaps between **intended data use** and **actual trust capability**.

It is not a certification.
It is a mirror.

How to use this assessment

This assessment should be completed collaboratively by representatives from:

- business leadership,
- data and analytics teams,
- legal / compliance,
- and domain experts.

Disagreement during this exercise is a signal of **missing alignment**, not failure.

Trust Readiness Dimensions

Evaluate each statement honestly.
If the answer is "it depends" or "usually," treat it as **No**.

1. Data Ownership & Accountability

☐ For every critical dataset, there is a clearly named owner
☐ Owners are empowered to approve or restrict data usage
☐ Accountability for data correctness is explicit, not implied
☐ Ownership survives organizational or system changes

If any box is unchecked:
Trust assertions are fragile and likely to fail under scrutiny.

2. Provenance & Context Awareness

☐ Data sources are known and documented
☐ Original purpose and collection context are preserved
☐ Contextual constraints (legal, contractual, situational) are visible
☐ Loss of context is treated as a trust-degrading event

If any box is unchecked:
Data may be accurate but not defensible.

© 2026 Martijn Wiggers
This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

3. Reliability Classification

☐ Data is classified by reliability, not just by sensitivity or quality
☐ Reliability levels are consistent across teams and domains
☐ Mixed-reliability datasets are explicitly recognized
☐ Reliability classifications are revisited over time

If any box is unchecked:
Trust decisions are implicit and inconsistent.

4. Decision Boundaries

☐ There are explicit rules linking reliability to allowed decisions
☐ Automation is gated by trust thresholds
☐ Human review is required where trust is insufficient
☐ Boundary violations are detectable and auditable

If any box is unchecked:
Automation may exceed what trust can justify.

5. Human Control Points

☐ Human intervention points are intentionally designed
☐ Overrides require justification and leave an audit trail
☐ Corrections trigger trust reassessment
☐ Responsibility for interventions is clearly assigned

If any box is unchecked:
Human judgment is informal and ungoverned.

6. Correction & Trust Recovery Handling

☐ Corrections preserve historical states
☐ The impact of corrections on past decisions is understood
☐ Trust recovery is conditional, not assumed
☐ Permanent trust degradation is accepted when appropriate

If any box is unchecked:
Trust is inflated and risk accumulates silently.

7. Organizational Culture

☐ “Defensible” is distinguished from “directionally correct”
☐ Assumptions are documented, not hidden
☐ Discomfort with uncertainty is tolerated
☐ Challenging data-driven decisions is culturally acceptable

If any box is unchecked:
Confidence may be mistaken for trust.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Interpreting the results

- **Mostly unchecked:**
  Your organization is not yet ready for trust-aware automation.
  Focus on governance and accountability before expanding AI usage.
- **Mixed results:**
  You likely operate with informal trust heuristics.
  Formalizing them will significantly reduce risk and friction.
- **Mostly checked:**
  You are structurally prepared to scale decision-making with trust as a constraint.

There is no “perfect score.”
Readiness is about **knowing your limits and respecting them**.

Final reflection

Ask one final question:

If a high-impact decision were challenged tomorrow, could we clearly explain why we trusted the data behind it?

If the answer is anything less than an unambiguous “yes,” trust readiness is incomplete.

Key takeaway

Trust readiness is not about perfection.
It is about **honesty, discipline, and accountability**.

Organizations that assess themselves rigorously gain not only safer systems, but **strategic advantage**: they know exactly where they can—and cannot—move fast.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Appendix A — Definitions & Glossary

This appendix defines the key terms used throughout this document.
All definitions are **normative** within the context of this framework and are intended to be used consistently.

Accountability

The explicit assignment of responsibility for the correctness, interpretation, and use of data or decisions derived from it.

Accountability is always human.
It cannot be delegated to systems, models, or processes.

Accuracy

The degree to which a data value corresponds to the actual real-world value it represents.

Accuracy is:

- quantitative,
- measurable,
- and independent of provenance or responsibility.

Accurate data may still be unreliable.

Automation

The execution of decisions or actions by systems without direct human involvement at the moment of execution.

Automation increases speed and scale, and therefore increases the **impact of trust failures**.

Confidence

A measure of certainty or probability assigned to a result, often produced by statistical or AI models.

Confidence reflects internal model certainty, **not** external defensibility or legitimacy.

Confidence must not be treated as a substitute for reliability.

Context

Information that defines the conditions under which data is valid, meaningful, and safe to use.

Context may include:

- legal or contractual constraints,
- exceptional circumstances,
- domain-specific meaning,
- temporal relevance,
- intent behind data creation.

Context modifies how reliability should be interpreted.

Correction

An intentional change to data intended to address an error, omission, or updated understanding.

Corrections may improve accuracy but do not automatically restore reliability.

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

Defensible (Defensibility)

The ability to justify the use of data or a decision under external scrutiny, including legal, regulatory, or ethical challenge.

Defensibility requires:

- known provenance,
- preserved context,
- and clear accountability.

Decision Boundary

A rule or constraint that defines which decisions are permitted at a given level of data reliability.

Decision boundaries translate trust into enforceable action limits.

Derived Data

Data produced through calculation, inference, aggregation, or modeling rather than direct observation or transaction.

Derived data is not a fact, even when deterministic.

Human-in-the-Loop

A system design pattern where humans explicitly intervene at defined control points to validate, approve, override, or correct outcomes.

Human-in-the-loop is a governance mechanism, not a quality fix.

Override

An intentional deviation from an automated or predefined outcome.

Overrides alter the trust state of data and must be explicit, justified, and auditable.

Provenance

Information describing the origin of data, including:

- source system or authority,
- method of collection,
- and transformation history.

Provenance is a prerequisite for reliability.

Reliability

The degree to which data can be responsibly relied upon for decision-making.

Reliability depends on:

- provenance,
- context,
- and accountability.

Reliability answers the question:

May we act on this data?

© 2026 Martijn Wiggers
This document provides conceptual guidance.
It does not constitute legal or regulatory advice.

All rights reserved.

Trust

The justified acceptance of data as a basis for action within defined boundaries.

Trust is not a feeling, score, or assumption.
It is an operational condition tied to responsibility and risk.

Trust Decay

The gradual loss of data reliability as data moves, transforms, or loses context and accountability.

Trust decay can occur even when data remains accurate.

Trust Recovery

The governed process by which data reliability may be partially or fully restored after correction or override.

Trust recovery is conditional and not guaranteed.

Trust Readiness

An organizational state in which structures, culture, and governance are sufficient to manage data trust explicitly and consistently.

Trust readiness is not a technical maturity level.

Quality (Data Quality)

The degree to which data is fit for a specific intended purpose.

Quality is context-dependent and does not imply reliability.

Key Clarification

- **Accuracy answers:** *Is this value correct?*
- **Quality answers:** *Is this data fit for purpose?*
- **Confidence answers:** *How certain is a model or person?*
- **Reliability answers:** *May we rely on this data for decisions?*

Only reliability determines defensible action.

Closing note

These definitions are intentionally strict.
Relaxing terminology may improve convenience, but it weakens trust.

Organizations that treat definitions loosely will inevitably treat decisions loosely as well.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.

About the Author

Martijn Wiggers operates in decision-critical domains where governance is under scrutiny and decisions must be defensible.
His work establishes practical frameworks for making data trust explicit, governable, and auditable in complex, data-driven or AI-driven environments.

© 2026 Martijn Wiggers

This document provides conceptual guidance.
It does not constitute legal or regulatory advice.
All rights reserved.