




What we mean by "We uphold NIST AI Risk Management principles"

April 2026

Our Commitment to Trustworthy AI

StoryDesk is built on the principles of the NIST AI Risk Management Framework (AI RMF) — one of the most widely recognised global frameworks for trustworthy AI. This document explains what that means, how the framework works, and how StoryDesk applies it in practice.

 This is not a certification. NIST AI RMF is a voluntary framework — there is no government body that certifies compliance the way a safety mark certifies an appliance. When we say we uphold these principles, we are making a public commitment to designing, building, and operating our platform in line with the framework's standards for trustworthy AI. It is a commitment we hold ourselves to, and one our users and partners can hold us to.

Part 1: What is the NIST AI Risk Management Framework?

The NIST AI Risk Management Framework was released by the National Institute of Standards and Technology, a United States federal agency that develops standards across science, technology, and industry. The AI RMF provides organisations with a structured approach to identifying, measuring, and managing the risks that come with deploying AI systems.

In practical terms, the framework represents a shift from "making AI work" to "making AI work safely." It addresses the gap between what AI systems can do and what they should be trusted to do — particularly when AI-generated content reaches a public audience.

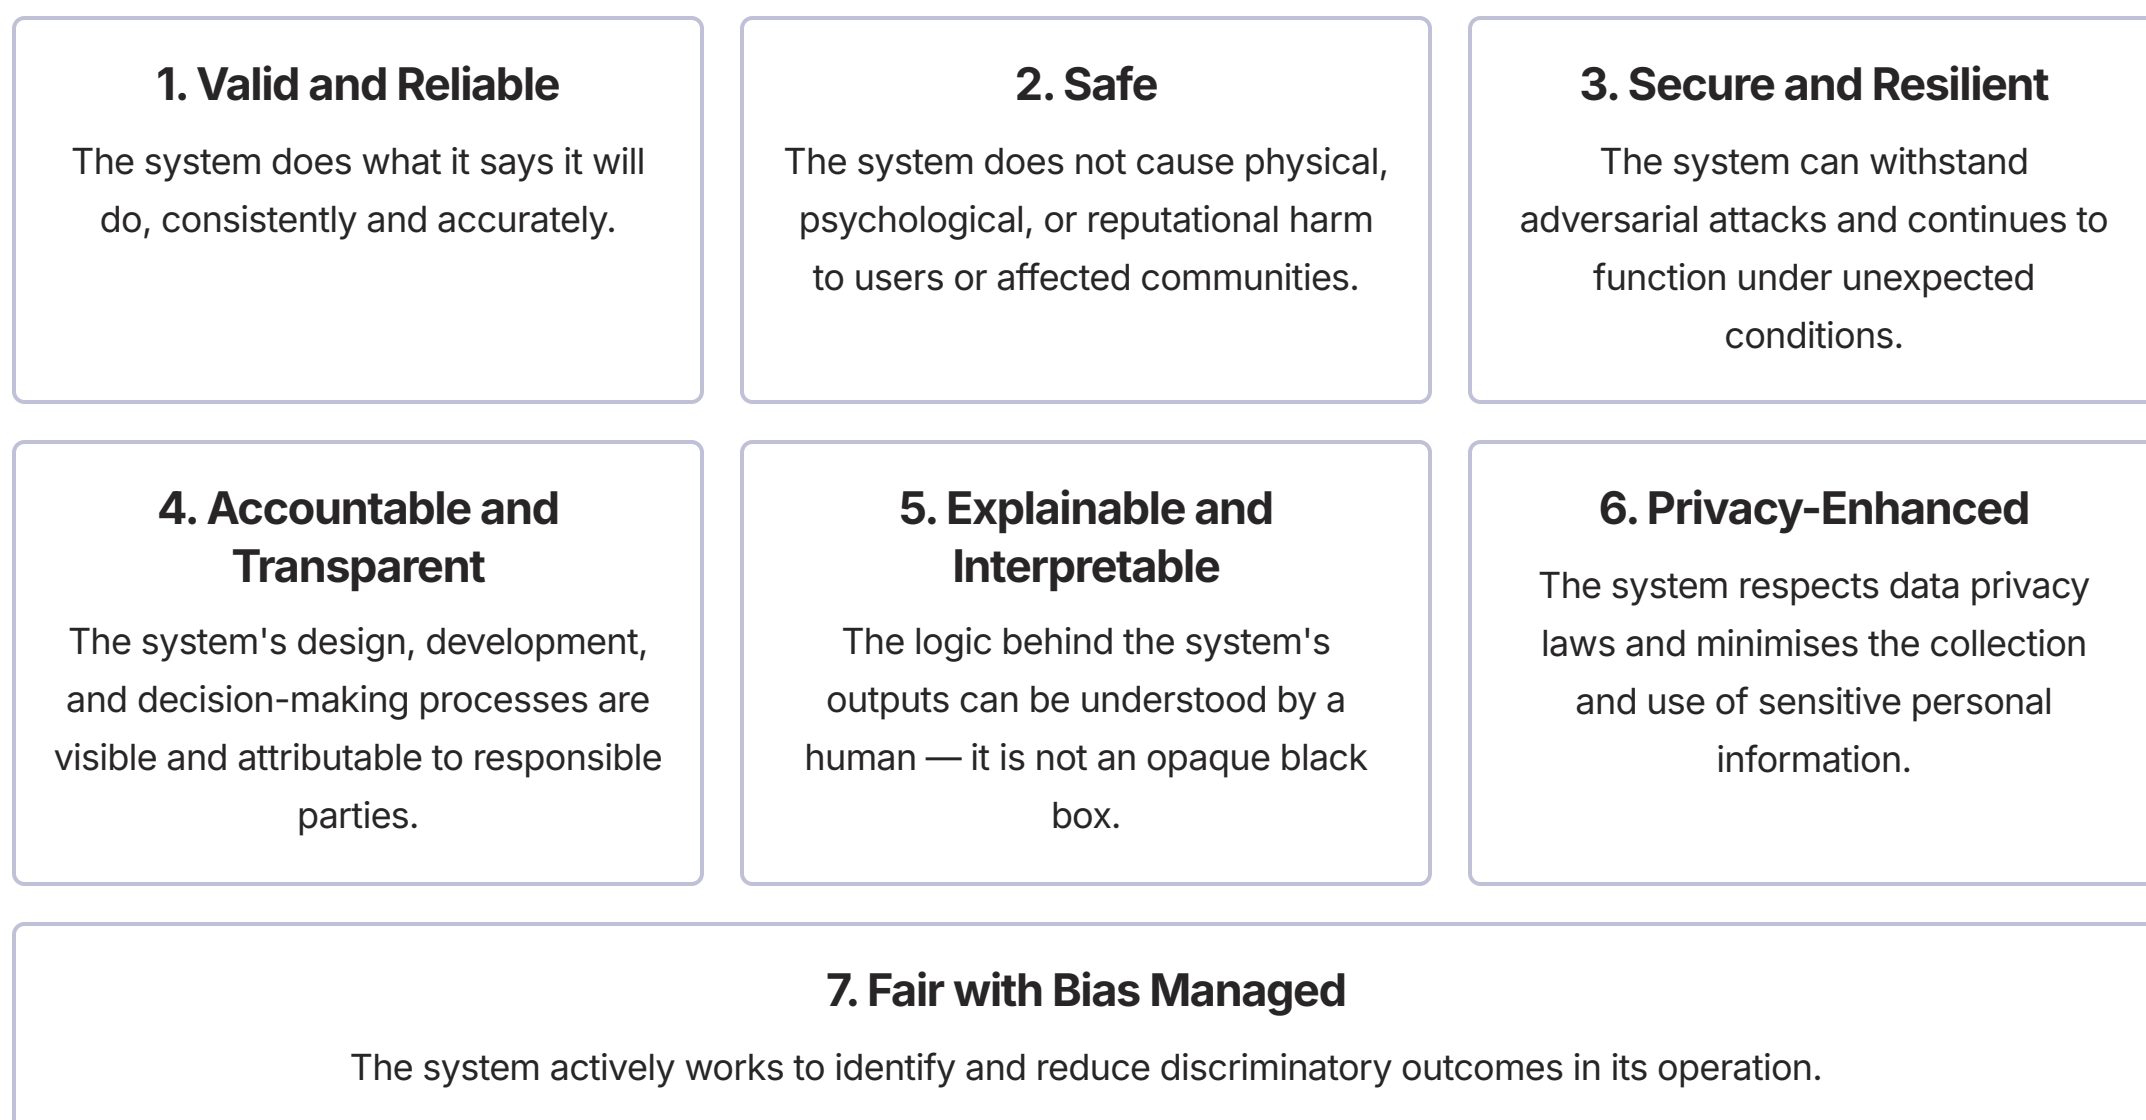
The four core functions

The framework is structured around four high-level functions. Together, they form a continuous cycle of responsible AI management:



The seven characteristics of trustworthy AI

The framework defines seven specific traits that a trustworthy AI system should demonstrate:



Part 2: How StoryDesk applies these principles

StoryDesk is an AI-powered content creation and publishing platform. Our users — marketing teams, agencies, journalists, independent creators, and regulated industries — use AI to research, draft, and publish content. The risks in this context are specific and measurable: AI-generated hallucinations, brand voice drift, unverified claims reaching a public audience, and content that fails to meet editorial or regulatory standards.

We address these risks through a three-stage governance pipeline that maps directly to the NIST framework's core functions.

MAP

Source Grounding

Before content is generated, StoryDesk maps the information landscape. Our unified search draws from eight source types — including news databases, PubMed, Google, Substack, and YouTube — and saves every source to a Source Library with full provenance tracking. The AI generates from material the creator has selected and can verify, not from the model's training data alone.

What this addresses:

- Reduces AI hallucination by anchoring content in retrievable sources
- Flags false attributions so creators can verify claims before publication

"Every claim traces back to a source the creator can verify."

MEASURE

Brand Safety Filtering

Before content reaches the creator for review, it passes through a filtering layer that measures the output against the organisation's editorial standards. Tone of Voice profiles enforce brand consistency per author. Banned word lists and AI vocabulary replacement catch off-brand language and the generic AI phrasing (AI slop) that signals machine-generated text to increasingly AI-literate audiences.

What this addresses:

- Prevents brand voice drift across content and contributors
- Eliminates off-brand language before human review and reduces regulatory risk

"Quality control at the content level"

MANAGE

Human-in-the-Loop Review

Every piece of AI-generated content enters a mandatory staging area before publication. The creator sees both the content and the underlying sources side by side, enabling a fully informed review decision. The creator approves, edits, or rejects each draft. No content bypasses this gate.

What this addresses:

- Maintains human accountability for every published piece
- Satisfies the NIST principle that high-stakes AI outputs require human oversight

"The AI drafts. The Creator decides. No unreviewed AI content leaves the platform"

Part 3: How this works in practice



From source to published post

Most AI content tools generate text from a prompt alone.

StoryDesk is designed to link every piece of content back to a verified source — a news article, a transcript, a research paper, or a curated entry in the creator's Source Library. The AI is tethered to reality, not generating unconstrained model output.



Multi-platform adaptation

StoryDesk publishes to multiple platforms. Each has unique conventions for tone, format, and length. The content intelligence pipeline applies platform-specific rules at the generation stage, so content is adapted to each channel's context.



The creator as decision-maker

The NIST framework emphasises that high-stakes AI systems should not operate on a "set it and forget it" basis. StoryDesk is built as a co-pilot, not an autopilot. The editorial staging area ensures that a human makes a conscious, accountable decision on every draft before it reaches an audience.

Why this matters

In a market full of AI content tools, most compete on speed and volume. **StoryDesk competes on trust.** Following NIST AI Risk Management principles means we are:



Proactive, not reactive

We test and screen content before it publishes, not after damage is done.



Accountable to a recognised standard

We are not grading our own homework. We measure our practices against a globally recognised framework developed by an independent standards body.



Transparent about what this means — and what it doesn't

NIST AI RMF is a voluntary framework. Following it is a public commitment, not a certification. We are clear about that distinction because transparency is itself a NIST principle.

We believe that trustworthy AI content tools will define the next generation of publishing platforms. StoryDesk is built to lead that standard, not follow it.