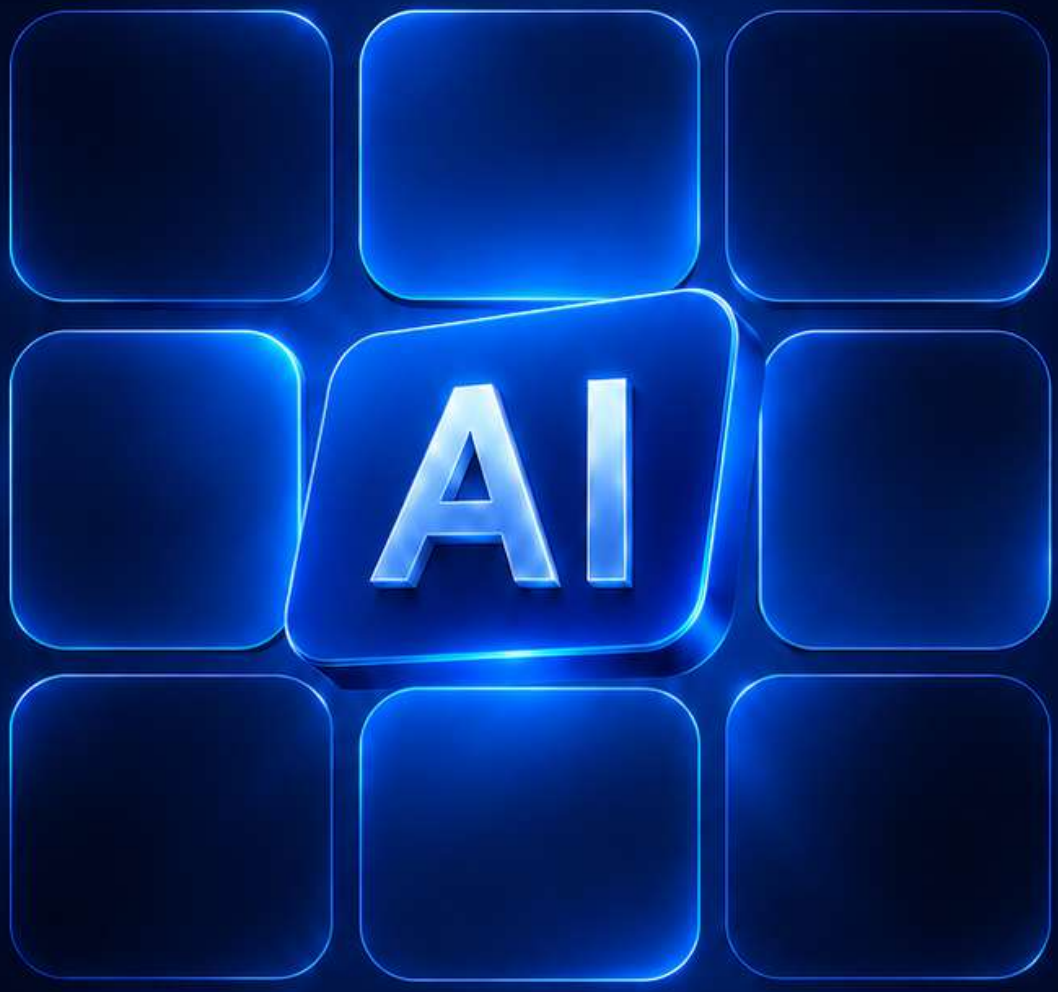
SAND STUDIO



SAND STUDIO
Responsible AI General Guide

FOR EXTERNAL RELEASE

2026

CONTENTS

- About this guide 03
- Bias Management Goals 04
- Transparency and Explainability Goals 05
- Accountability Goals 07
- Data Rights Goals 09
- Documentation and Monitoring Goals 11
- Safety and Reliability Goals 13
- Security and Robustness Goals 15
- Notes 16

About this Guide



As AI becomes increasingly integrated into products, services, and business processes, AI-related activities become more diverse and span different stages of the AI lifecycle. SAND STUDIO has established management systems that provide an important foundation for applying responsible AI principles in practice. As AI use cases continue to expand, these existing measures need to be supported by practical requirements to address AI-related risks throughout the AI lifecycle.

Building on this foundation, this Guide describes SAND STUDIO's approach to responsible AI governance and provides a framework for considering AI-related risks across the development, deployment, and use of AI systems. It is intended to support product, technical, and governance decisions while giving customers, partners, users, and other stakeholders greater visibility into how we govern AI and manage AI-related risks.

Responsible AI practices will continue to evolve alongside developments in AI technology and research, as well as changes in laws, business uses, and user expectations. Progress depends on collaboration among industry, researchers, public institutions, civil society, customers, and users. We will continue to learn from these developments and share our experience to help make responsible AI practices clearer and more practical.

Bias Management Goals

Goal B1: Fair and equitable AI outcomes

Where an AI system or use case may present a material fairness risk, SAND STUDIO assesses whether it could produce unfair outcomes and defines appropriate fairness objectives.

Requirements
B1.1: Identify stakeholders, use cases, and groups that may be affected by AI system outputs or decisions. Consider demographic, social, cultural, geographic, and other characteristics relevant to the intended use.
B1.2: When a material fairness risk is identified, define and document fairness objectives and evaluation criteria suited to the intended use and potential impact of the AI system.
B1.3: Use design, data, product, or operational measures intended to reduce the risk of discriminatory or unjustifiably disparate outcomes.
B1.4: When fairness objectives have been defined, evaluate relevant outputs or decisions against them. Document material disparities and determine whether further controls, use limitations, or another documented response are needed.

Goal B2: Bias identification and mitigation

SAND STUDIO examines relevant sources of bias and takes steps to reduce material bias risks.

Requirements
B2.1: Identify potential sources of bias in the system's data, design, development processes, outputs, and human interactions.
B2.2: When relevant and legally permitted, assess the representation of identified groups in datasets used to develop, train, validate, or test the AI system.
B2.3: Test for material differences in performance or outcomes across relevant groups and use conditions when such testing is feasible and justified by risk.
B2.4: Define and document how identified bias risks will be assessed and how the effectiveness of mitigation measures will be validated before deployment. Reassess the need for validation when material changes to the model, data, intended use, or deployment context could affect those risks.
B2.5: Implement measures intended to reduce material bias risks and complete any validation required under B2.4 at the lifecycle stage appropriate to the system and its risk.

Transparency and Explainability Goals

Goal T1: AI intelligibility and explainability

SAND STUDIO provides information intended to help users and other stakeholders understand relevant AI systems and use their outputs responsibly.

Requirements
T1.1: Identify users and other stakeholders who need to understand the AI system, its outputs, or the decisions it supports.
T1.2: Determine what those stakeholders need to know about the system's purpose, operation, capabilities, limitations, uncertainty, and potential impact. Provide that information through the system, user experience, documentation, or other suitable means.
T1.3: When evaluation is needed, define and document the methods and criteria used to assess the effectiveness of the selected transparency and explainability measures.
T1.4: Review the selected measures after material changes to the system, intended use, deployment context, or stakeholder needs.

Goal T2: Third-party AI transparency review

SAND STUDIO reviews third-party AI systems when their proposed use could introduce material risk.

Requirements
T2.1: Before integration or use, identify third-party AI systems that may introduce material risks when integrated into or used by SAND STUDIO products, services, or internal processes.
T2.2: Review information reasonably available from the third-party provider. This may include intended uses, capabilities, limitations, performance, data practices, security information, and known risks.
T2.3: Assess whether the available information is sufficient to understand the material risks relevant to the proposed use. Consider relevant dependencies and determine whether additional controls, restrictions, contingency or exit arrangements, or further review are needed.
T2.4: Document the review result, material information gaps, limitations, conditions of use, and SAND STUDIO's response to the identified risks. Reassess the third-party AI system when material changes to the provider, model, intended use, or risk make further review appropriate.

Goal T3: Clear disclosure of AI interaction

SAND STUDIO discloses AI involvement when required by the business context, potential impact, applicable obligations, or reasonable user expectations.

Requirements
T3.1: Identify relevant product entry points, workflows, user interactions, and AI-generated content scenarios in which users could reasonably believe they are interacting with a person or consuming human-generated content.
T3.2: When disclosure is required, present it in a form, location, and timing suited to the intended users and context.
T3.3: Define and document disclosure requirements and relevant exceptions for the scenarios identified under T3.1, including cases where AI involvement is already clear from the context.
T3.4: When needed, evaluate whether the disclosure is noticeable, understandable, and accessible to the intended users.

Accountability Goals

Goal A1: Clear accountability and responsibility

SAND STUDIO assigns ownership and decision-making responsibilities for AI systems according to their lifecycle stage and risk.

Requirements
A1.1: Assign an accountable owner to each AI system within scope and identify the stakeholders responsible for relevant activities across its design, development, acquisition, deployment, operation, oversight, governance, and retirement.
A1.2: Define and document their roles, responsibilities, decision rights, handoffs, and escalation paths at the level needed for the system. As applicable, these responsibilities include risk and impact assessments, model and service change management, review of templates before release, and ongoing improvement.
A1.3: Determine who is responsible for reviewing, validating, approving, or confirming AI-generated outputs or decisions when human review is required.
A1.4: Keep records of material ownership, approval, escalation, and risk-management decisions in accordance with applicable recordkeeping requirements.

Goal A2: Risk-based human oversight and control

SAND STUDIO determines the level of human oversight required for an AI system based on its intended use and potential impact.

Requirements
A2.1: Identify AI systems and use cases that may produce significant or far-reaching effects on individuals or groups. Determine the level and form of human oversight needed.
A2.2: Define and document the conditions in which human review, approval, intervention, override, correction, or monitoring is required. For high-impact, sensitive, or difficult-to-reverse automated actions, document any required human review, conditions for stopping or escalating the workflow, available recovery or rollback measures, and logging requirements in the relevant product or workflow documentation.
A2.3: When human intervention is required, identify the authorized personnel and provide them with the information, authority, access, and available mechanisms needed to review the output or take the defined action.
A2.4: Establish escalation paths for unexpected, ambiguous, harmful, or otherwise material system behavior.
A2.5: Before deploying an AI system that may make or support consequential decisions affecting individuals or groups, determine and document any required human oversight, privacy safeguards, legal review, and approval.

Goal A3: Human-AI collaboration and autonomy

SAND STUDIO defines the intended roles of people and AI systems and considers how AI may affect human judgment and user choice.

Requirements
A3.1: Define and document the intended roles and boundaries of people and the AI system based on intended use and risk.
A3.2: Identify situations in which human judgment, intervention, or decision-making should take precedence. When relevant, design a process for handing control from the AI system to an authorized person.
A3.3: Provide users with information and controls proportionate to the consequences of relying on AI-generated outputs or recommendations.
A3.4: Assess the risk of over-reliance, automation bias, or loss of meaningful user choice when those risks could result in material harm. Use safeguards proportionate to the identified risk.

Data Rights Goals

Goal D1: Data governance and rights

SAND STUDIO identifies and manages rights, restrictions, and responsibilities associated with data used or generated by AI systems.

Requirements
D1.1: Identify applicable data rights, licenses, usage restrictions, information barriers, and contractual obligations relevant to data used by the AI system. When relevant, this includes customer inputs, knowledge-base content, data retrieved through RAG, data exchanged through tools or APIs, and data or content obtained from third-party models or providers.
D1.2: Define and document permitted data uses across relevant lifecycle stages. These may include system development, configuration, deployment, operation, evaluation, and improvement.
D1.3: Assign ownership and operational responsibility for data used or generated by the AI system, including responsibility for maintaining relevant source and authorization records and managing applicable third-party intellectual property requirements.
D1.4: Use system and data-flow controls intended to enforce applicable information barriers, usage restrictions, and contractual obligations.

Goal D2: Data quality and suitability

SAND STUDIO assesses whether the data is suitable for the AI system's intended purpose and risk profile.

Requirements
D2.1: Define and document data requirements based on the system's intended purpose, expected use conditions, evaluation needs, and risk profile. Where relevant, these requirements apply to knowledge-base content, data used for RAG, and data used to test or evaluate the system.
D2.2: Assess relevant characteristics of available data against the requirements defined in D2.1. These characteristics may include quality, quantity, relevance, completeness, representativeness, provenance, and timeliness.
D2.3: Determine whether available data sources satisfy the documented requirements and can be maintained, updated, or removed as appropriate for the intended use and applicable retention requirements. When they do not, determine whether remediation, additional testing, restricted use, or another control is needed.
D2.4: Identify and document material data limitations that may affect performance, safety, reliability, fairness, or intended outcomes. Determine how to treat or communicate those limitations. Reassess data suitability when material data changes or data drift could affect the system's performance or risk profile.

Goal D3: Data protection and access control

SAND STUDIO selects and applies controls to protect AI inputs and outputs based on data sensitivity, risk, and applicable obligations.

Requirements
D3.1: Identify the privacy, security, confidentiality, legal, and contractual requirements that apply to relevant AI inputs and outputs.
D3.2: Define and document protection requirements based on data sensitivity, intended use, and applicable obligations. Where appropriate, use existing information security management processes to implement these requirements.
D3.3: Use controls intended to reduce the risk of unauthorized access, disclosure, alteration, loss, inappropriate retention or deletion, or misuse of AI inputs and outputs.
D3.4: Depending on data sensitivity and risk, controls may include authentication, authorization, data masking, encryption where appropriate, logging, monitoring, retention and deletion controls, access revocation, and the preservation of relevant records for investigations or legal obligations.
D3.5: Review access to sensitive AI data and outputs at a frequency based on risk and applicable policy. Revoke access when it is no longer needed, including after account or service termination where applicable. Investigate and address unauthorized or inappropriate access through the relevant SAND STUDIO process.

Documentation and Monitoring Goals

Goal M1: AI system inventory

SAND STUDIO maintains an inventory of AI systems that fall within the scope of its AI management processes.

Requirements
M1.1: Identify AI systems developed, provided, acquired, integrated, or deployed by SAND STUDIO that meet the applicable inventory criteria, and record their intended uses.
M1.2: Maintain a documented inventory of applicable AI systems, including their deployment context, relevant risk classification, accountable owner, and lifecycle status. Where relevant, record key third-party providers or dependencies needed to understand or manage the system.
M1.3: Update the inventory when an AI system is introduced, materially changed, repurposed, suspended, resumed, or retired.

Goal M2: AI lifecycle documentation

SAND STUDIO maintains lifecycle documentation proportionate to the risk, complexity, lifecycle stage, and governance needs of an AI system.

Requirements
M2.1: Identify and document relevant stakeholders, roles, intended uses, material design decisions, development or acquisition processes, and lifecycle stages.
M2.2: When applicable and reasonably available, document the datasets used to develop, train, validate, test, improve, or evaluate the AI system, including their purpose and known limitations.
M2.3: Record material decisions, events, changes, approvals, and other information required by SAND STUDIO's governance processes throughout the system lifecycle.
M2.4: Maintain technical and governance records in a form that supports the level of traceability, review, maintenance, and audit needed for the system.
M2.5: When an AI system is retired, document the retirement decision, update the AI system inventory, revoke access that is no longer needed, and address affected third-party services or dependencies where applicable. Retain, archive, or delete relevant system components, data, and records under SAND STUDIO's control in accordance with applicable retention, legal, and contractual requirements.

Goal M3: AI risk and impact assessment

SAND STUDIO uses documented criteria to determine when an AI risk or impact assessment is required.

Requirements
M3.1: For a new AI system or intended use, determine whether an impact assessment is required based on its potential risks and effects.
M3.2: Complete any required assessment before the relevant approval, provision, or deployment decision for the AI system.
M3.3: Document the assessment results, identified risks and impacts, selected controls, accountable owners, and required follow-up actions. Use the results to inform risk treatment and approval decisions.
M3.4: Review whether an assessment needs to be updated following a material change to the system, intended use, deployment context, provider, data, or risk profile. Update the assessment when the change could materially affect the system's risks or impacts.

Goal M4: Ongoing monitoring and review

SAND STUDIO defines monitoring and review activities for deployed AI systems according to their risk and operating context.

Requirements
M4.1: Define and document the monitoring objectives, methods, thresholds, owners, and review frequency needed for the system.
M4.2: Monitor the system at a level suited to its risk to identify performance issues, unexpected outcomes, emerging risks, control failures, and material impacts.
M4.3: Record material monitoring results, identified issues, corrective actions, improvements, and significant lifecycle events.
M4.4: Conduct scheduled or event-driven reviews at a frequency based on risk, applicable policy, contractual commitments, and operational needs. Include reviews following material changes, incidents, significant performance degradation, or changes in applicable obligations when needed.
M4.5: Investigate material issues and track corrective actions through completion, including an effectiveness review when needed.

Safety and Reliability Goals

Goal S1: Reliable and safe operation

SAND STUDIO defines and evaluates expected operating conditions, performance criteria, and safety controls for AI systems.

Requirements
S1.1: Define and document the intended use, expected operating conditions, performance expectations, and relevant foreseeable use conditions.
S1.2: When measurable criteria are useful, define performance thresholds, acceptable error levels, or other release and operating criteria proportionate to the intended use and risk.
S1.3: Design and evaluate the system against the defined operating conditions and performance criteria selected for the intended use.
S1.4: Identify material conditions in which the system may not operate as intended. Determine how those limitations should be controlled or communicated to relevant stakeholders.

Goal S2: Failure management and misuse resilience

SAND STUDIO identifies foreseeable failure and misuse scenarios and uses controls proportionate to their potential impact.

Requirements
S2.1: Identify and document relevant failures, misuse scenarios, manipulation attempts, and adverse operating conditions. Where relevant to the system, these may include significant model or data drift, harmful or materially biased outputs, materially inaccurate or fabricated outputs, unexpected automated behavior, and material control failures.
S2.2: Assess the potential effects of the scenarios identified in S2.1 on users, other stakeholders, dependent systems, and operations.
S2.3: Use measures intended to prevent, detect, limit, or recover from the identified failures and misuse scenarios.
S2.4: When needed, establish fallback, recovery, rollback, feature restriction, or contingency measures for failures that cannot be sufficiently addressed through system design.
S2.5: Define and document conditions under which AI outputs or actions should be restricted, interrupted, reviewed, or escalated for human intervention. Record and address material events through the applicable incident-management process. Based on risk, such handling may include escalation, investigation, corrective action, and effectiveness reviews.

Goal S3: Testing and validation

SAND STUDIO defines testing and validation activities based on intended use, risk, operating conditions, and material system changes.

Requirements
S3.1: Define and document the test objectives, methods, data, environments, acceptance criteria, and reviewers appropriate to the system.
S3.2: Conduct testing and validation before deployment and, when warranted by risk, after material changes, at the stages appropriate to the system.
S3.3: Select test conditions proportionate to risk. These may include normal use, foreseeable variation, relevant edge cases, failure scenarios, adverse conditions, and misuse scenarios. Where relevant, feasible, and legally permitted, additional testing may address bias or fairness, data or model drift, prompt injection, RAG retrieval and grounding, tool or API interactions, and recovery or rollback.
S3.4: Compare results with the defined criteria. Document material gaps and determine whether remediation, restricted use, another control, or a documented risk decision is needed before release or continued use.

Security and Robustness Goals

Goal R1: AI security and resilience

SAND STUDIO uses risk-based security controls intended to reduce the likelihood and impact of threats to AI systems and associated data.

Requirements
R1.1: Identify relevant threats and vulnerabilities affecting the system, its models, data, interfaces, infrastructure, dependencies, providers, and users. Where relevant to the system, these may include prompt injection, data poisoning, model extraction, and risks arising from third-party APIs or dependencies.
R1.2: Define and document security requirements based on the threat assessment, system architecture, intended use, data sensitivity, and risk.
R1.3: Use controls intended to reduce the risk of unauthorized access, disclosure, modification, disruption, manipulation, or misuse.
R1.4: Apply secure-by-design practices throughout the AI system lifecycle at a level suited to risk and applicable policy. These practices may include: • Identity and access management • Data protection • Secure development • Dependency management • Configuration management • Logging and monitoring • Security testing • Vulnerability management
R1.5: Apply SAND STUDIO's incident-management processes to security events that may compromise the confidentiality, integrity, availability, functionality, or intended behavior of the AI system.

Goal R2: System compatibility and operational resilience

SAND STUDIO assesses and manages the risk that an AI deployment may disrupt dependent systems, services, or operational processes.

Requirements
R2.1: Identify material dependencies, interfaces, integration requirements, capacity needs, and operational constraints relevant to the AI system and deployment environment.
R2.2: Assess potential effects of deployment or operation on existing systems, processes, services, users, and resources.
R2.3: Define the interoperability and compatibility requirements needed for the intended deployment.
R2.4: Validate, at a level suited to the system's risk, that deployment satisfies the selected performance, availability, compatibility, and service quality criteria.
R2.5: When needed, establish safeguards, fallback measures, capacity controls, or operational restrictions intended to prevent or limit disruption to dependent systems and services.

Notes

Scope and application

This Guide describes SAND STUDIO's general approach to responsible AI governance. The applicability and implementation of each Goal and Requirement may vary depending on the AI system's intended use, risk profile, lifecycle stage, deployment context, and applicable legal or contractual requirements.

This Guide is provided for general information. It does not create or modify any warranty, representation, or contractual commitment, nor does it guarantee the absence of risks or any particular outcome in connection with AI-related activities.

Management systems and certifications

This Guide also supports SAND STUDIO's AI and information security management systems, which are certified to ISO/IEC 42001:2023 and ISO/IEC 27001:2022, respectively.

ISO/IEC 42001:2023

The certification scope covers the design, development, and provision of the GoInsight.AI enterprise AI collaboration and automation platform, together with related support using third-party foundation models.

ISO/IEC 27001:2022

The certification scope covers the provision of an endpoint management solution and an AI automation platform.

These certifications relate to the management systems within those scopes and do not constitute certification of any individual product, AI system, customer deployment, or AI output, or guarantee any particular performance or outcome of such products, systems, deployments, or outputs.

Update to this Guide

This Guide reflects SAND STUDIO's current approach to responsible AI governance. SAND STUDIO may update the Guide from time to time to reflect development in technology, changes in law, stakeholder expectations, or its management practices.