

SELECTED NIST LINKS

Selected connections to NIST guidance

AI Content Verification Toolkit for Small Publishers

Publisher: ArizonaAI.com | Project lead: Cenk Brown | Release 1.1.0

This document links parts of the toolkit to AI RMF 1.0 and the Generative AI Profile. It does not claim full coverage, equal scope, compliance, use by an organization, or NIST endorsement.

Read each row’s limit

A form can help with a task without proving anyone did it. A source check confirms its ID and meaning. Check the toolkit evidence and its dated review status separately.

Scope label	What it means in this mapping
Practical support for a selected aspect	14 rows name a limited task the toolkit helps with.
Contextual connection	2 rows explain a related idea. They do not claim the outcome was reached.
Not covered	2 rows name tasks the toolkit does not cover.

Source baseline

SRC05 | Artificial Intelligence Risk Management Framework
NIST AI 100-1, AI RMF 1.0, January 2023. Each row gives printed and one-based PDF page locators.

SRC06 | Generative Artificial Intelligence Profile
NIST AI 600-1, Generative AI Profile, July 2024. Each selected action retains its exact identifier.

HTML mapping | Source register | Validation record

No row is marked pending in the source data. Each keeps its dated review status. This is not a new NIST review.

This release has not been sent to or accepted by NIST. Our original guide uses CC BY 4.0. Read LICENSE-CONTENT.md for its scope and exclusions. Third-party documents keep their own rights. You can read the HTML mapping or use PDF bookmarks for each row. We do not claim PDF/UA certification.

GOVERN 1.5

Scope: Practical support for a selected aspect

Plan ongoing checks. Name an owner and choose how often to review.

SRC05 | Artificial Intelligence Risk Management Framework

NIST AI 100-1 | AI RMF 1.0 | Section 5; Table 1 | Printed page 23; PDF page 28 (one-based).

Toolkit location: [guide.html#monitoring](#)

Toolkit evidence: The guide asks for a check date or trigger, a reason, and an owner. The schedule form records the work to check, owner and backup, priority, selection method, findings, and next change.

Limit: The publisher must assign people and do the checks. A blank form does not show that this happens.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

GOVERN 2.1

Scope: Practical support for a selected aspect

Record who handles AI risks and how people should contact them.

SRC05 | Artificial Intelligence Risk Management Framework

NIST AI 100-1 | AI RMF 1.0 | Section 5; Table 1 | Printed page 23; PDF page 28 (one-based).

Toolkit location: [guide.html#human-review](#)

Toolkit evidence: The guide separates checking, human approval, a second review, and asking for help. The roles form asks for a contact route, correction owner, backup, and a reason to check the roles again.

Limit: A small-team form does not show that the whole organization has assigned or shared these roles.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

MAP 1.1

Scope: Practical support for a selected aspect

Record the intended use, setting, users, effects, and limits.

SRC05 | Artificial Intelligence Risk Management Framework

NIST AI 100-1 | AI RMF 1.0 | Section 5; Table 2 | Printed page 26; PDF page 31 (one-based).

Toolkit location: [guide.html#review-context](#)

Toolkit evidence: The guide asks who will use the content, what could happen if it is wrong, what time it describes, and what was not checked. Set that context before judging claims.

Limit: This covers the chosen publishing task. It does not fully assess the wider legal, social, environmental, business, or life-cycle setting.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

MAP 3.5

Scope: Practical support for a selected aspect

Define, assess, and record how people oversee the work.

SRC05 | Artificial Intelligence Risk Management Framework

NIST AI 100-1 | AI RMF 1.0 | Section 5; Table 2 | Printed page 27; PDF page 32 (one-based).

Toolkit location: [guide.html#decision-rules](#)

Toolkit evidence: The guide names human roles, separates Ready from Approved, blocks unresolved issues, and ties approval to one version. Follow-up questions ask whether owners acted and whether important changes were missed.

Limit: A recorded decision does not prove that oversight works or that an organization uses the process.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

MEASURE 1.1

Scope: Practical support for a selected aspect

Choose which risks to measure first. Record risks you did not measure.

SRC05 | Artificial Intelligence Risk Management Framework

NIST AI 100-1 | AI RMF 1.0 | Section 5; Table 3 | Printed page 29; PDF page 34 (one-based).

Toolkit location: [guide.html#metrics](#)

Toolkit evidence: The guide starts with risks that matter to readers. Four measures ask for dates, totals, selection rules, exclusions, and gaps. The form records priority, what is counted, how, what is not measured, and the next action. Results stay blank until measured.

Limit: These are example measures for publishing, with a blank record for choosing measures and noting gaps. The user must choose suitable measures and gather data. This does not assess all AI risks or prove results.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

MEASURE 2.1

Scope: Contextual connection

Record the data, measures, and tools used in testing.

SRC05 | Artificial Intelligence Risk Management Framework

NIST AI 100-1 | AI RMF 1.0 | Section 5; Table 3 | Printed page 29; PDF page 34 (one-based).

Toolkit location: [methods.html#evidence-classes](#)

Toolkit evidence: The methods page keeps practice cases separate from software tests. The guide asks for the version, task, materials, steps, kind of reviewer, findings, measures, and limits. These show how to record a test. They are not AI-system testing, evaluation, verification, and validation (TEVV) results.

Limit: Made-up cases and software checks do not test an AI system or show real-world AI accuracy.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

MEASURE 2.5

Scope: Contextual connection

Show whether results are valid and reliable. Record where they may not apply.

SRC05 | Artificial Intelligence Risk Management Framework

NIST AI 100-1 | AI RMF 1.0 | Section 5; Table 3 | Printed page 29; PDF page 34 (one-based).

Toolkit location: [guide.html#evaluation](#)

Toolkit evidence: The guide separates made-up practice, dated source examples, software checks, agent checks, and human studies. It says these examples do not prove broad reliability or results for an AI system.

Limit: The toolkit does not prove that a model is valid or reliable, or that its results apply elsewhere. It helps explain why a broad claim may lack proof.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

MEASURE 2.8

Scope: Practical support for a selected aspect

Check and record risks from unclear information or responsibility.

SRC05 | Artificial Intelligence Risk Management Framework

NIST AI 100-1 | AI RMF 1.0 | Section 5; Table 3 | Printed page 30; PDF page 35 (one-based).

Toolkit location: [guide.html#transparency](#)

Toolkit evidence: The guide asks whether readers could wrongly assume independent reporting, a current check, or a wider review. It also asks who owns corrections. Record the risk and the source credit, date, limit, or contact detail used to address it.

Limit: The questions address selected risks about credit, dates, ownership, and disclosure in publishing. They do not fully assess an AI system's transparency or accountability.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

MANAGE 4.1

Scope: Practical support for a selected aspect

Put monitoring, feedback, error response, recovery, and change steps in place.

SRC05 | Artificial Intelligence Risk Management Framework

NIST AI 100-1 | AI RMF 1.0 | Section 5; Table 4 | Printed page 33; PDF page 38 (one-based).

Toolkit location: [guide.html#monitoring](#)

Toolkit evidence: The guide names check triggers, an owner, and a correction route that the publisher must monitor. It asks about late or missed checks. Correction steps cover limiting spread, checking evidence and related content, choosing who to tell, and recording actions actually done.

Limit: This helps with a limited publishing task. It is not a full management program for an AI system in use.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

MANAGE 4.3

Scope: Practical support for a selected aspect

Tell the right people about incidents. Record the response and recovery.

SRC05 | Artificial Intelligence Risk Management Framework

NIST AI 100-1 | AI RMF 1.0 | Section 5; Table 4 | Printed page 33; PDF page 38 (one-based).

Toolkit location: [guide.html#corrections](#)

Toolkit evidence: The correction form records the content and version, evidence, owner, steps to limit spread, corrected words, new check, related content, who to tell, and where and when work finished. A plan to send a notice is kept separate from a sent notice.

Limit: The organization must decide who is affected, act, and record what it actually told people and fixed.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

GV-1.5-001

Scope: Practical support for a selected aspect

Assign people to check source history and watch for incidents on a schedule.

[SRC06 | Generative Artificial Intelligence Profile](#)

NIST AI 600-1 | Generative AI Profile | Section 3; GOVERN 1.5 | Printed page 16; PDF page 20 (one-based).

Toolkit location: [guide.html#monitoring](#)

Toolkit evidence: Source records include publishers, places to find evidence, access limits, and shared origins. The schedule assigns owners for changed sources, reader reports, and open issues. It records when to check and what was found.

Limit: An owner field and schedule are just a form. They do not prove that people followed through.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

GV-1.5-003

Scope: Practical support for a selected aspect

Keep a policy for retaining test records and content history.

[SRC06 | Generative Artificial Intelligence Profile](#)

NIST AI 600-1 | Generative AI Profile | Section 3; GOVERN 1.5 | Printed page 16; PDF page 20 (one-based).

Toolkit location: [guide.html#retention](#)

Toolkit evidence: The guide and policy form ask for an owner, record types including test records, a period or deletion trigger, reasons, permitted evidence, file links, access, deletion decisions, and policy review. They set no universal period.

Limit: Saving a file is not a retention policy. Each organization must choose how long to keep records. NIST does not set one period here for everyone.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

MS-2.5-001

Scope: Practical support for a selected aspect

Do not make broad capability claims from a small or informal test.

[SRC06 | Generative Artificial Intelligence Profile](#)

NIST AI 600-1 | Generative AI Profile | Section 3; MEASURE 2.5 | Printed page 31; PDF page 35 (one-based).

Toolkit location: [guide.html#evaluation](#)

Toolkit evidence: The guide keeps practice, source checks, software tests, and agent checks separate from human studies. It says not to infer detection rates, time savings, less harm, or model reliability from examples. Claims about results need the test steps, observations, and limits.

Limit: A small set of made-up cases cannot prove better results, less harm, or real-world AI accuracy.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

MS-2.5-003

Scope: Practical support for a selected aspect

Check output sources and citations before use and during later monitoring.

[SRC06 | Generative Artificial Intelligence Profile](#)

NIST AI 600-1 | Generative AI Profile | Section 3; MEASURE 2.5 | Printed page 31; PDF page 35 (one-based).

Toolkit location: [guide.html#source-evidence](#)

Toolkit evidence: The guide asks you to open sources, match identity and context, and record where the evidence is and what was checked. Each key claim links to evidence. Support, source suitability, and date fit stay separate. Later checks revisit changing sources and linked claims.

Limit: A person supplies the judgments. A source link alone cannot prove support, source fit, accuracy, or current truth.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

MS-2.8-003

Scope: Not covered

Keep traceable content history with methods that protect it from tampering.

[SRC06 | Generative Artificial Intelligence Profile](#)

NIST AI 600-1 | Generative AI Profile | Section 3; MEASURE 2.8 | Printed page 34; PDF page 38 (one-based).

Toolkit location: [methods.html#data-handling](#)

Toolkit evidence: The methods page says editable records, saved versions, and file hashes cannot prove identity, prove event times, or prevent tampering. The toolkit has no protected source-history system. This is why the action is Not covered.

Limit: Not covered. Editable files and hashes alone do not provide the protection this action describes.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

MS-2.8-004

Scope: Not covered

Test whether instructions for users of a generative AI system are adequate.

[SRC06 | Generative Artificial Intelligence Profile](#)

NIST AI 600-1 | Generative AI Profile | Section 3; MEASURE 2.8 | Printed page 34; PDF page 38 (one-based).

Toolkit location: [methods.html#validation](#)

Toolkit evidence: The methods page separates worksheet usability checks from tests of a generative AI system's user instructions. No such instruction test is reported. The guide also marks this action Not covered.

Limit: Not covered. This release has not tested those instructions with users. An AI walkthrough is not human research.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

MG-4.1-002

Scope: Practical support for a selected aspect

Keep and assess ways to monitor generative AI after it is put into use.

SRC06 | [Generative Artificial Intelligence Profile](#)

NIST AI 600-1 | Generative AI Profile | Section 3; MANAGE 4.1 | Printed page 44; PDF page 48 (one-based).

Toolkit location: [guide.html#monitoring](#)

Toolkit evidence: The guide asks whether source changes were found before readers used old information, what was missed, how it was found, and whether owners acted. The schedule form records evidence and planned changes to timing, selection, sources, or who to ask for help.

Limit: This only supports checks for made-up facts in published writing. It does not assess how well monitoring works or wider chemical, biological, radiological, nuclear (CBRN), or cyber risks.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

MG-4.3-002

Scope: Practical support for a selected aspect

Keep steps for tracking errors, close calls, and harmful effects.

SRC06 | [Generative Artificial Intelligence Profile](#)

NIST AI 600-1 | Generative AI Profile | Section 3; MANAGE 4.3 | Printed page 45; PDF page 49 (one-based).

Toolkit location: [guide.html#corrections](#)

Toolkit evidence: The guide and form keep reported errors, errors caught before publication, and observed harm separate. They record affected content, evidence, unknowns, owner, actions, a new check, and completion. A draft error is not treated as proof of harm.

Limit: A log does not prove a good response or that harm occurred. Actions and effects need evidence.

Recorded source check: October 6, 2026 at 11:51:04 PM Arizona time; Identifier, locator, and paraphrase checked

Recorded mapping review: Source and artifact reviewed

Before reusing this map, open the toolkit link and official passage. Check any changed document, page number, toolkit section, or scope. A link to one part does not mean the whole guidance is put into practice.