// ART-HLC-022 · PUBLISHED 8 NOVEMBER 2026

NIST and Trustworthy AI: The Lab That Never Fines Anyone

The National Institute of Standards and Technology sets standards nobody has to follow, and its AI framework is becoming the one that laws, buyers and accreditors point to, in healthcare first.

Jason Newell · MAX Research Collective · Article ART-HLC-022 · Published 8 November 2026
Licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0), https://creativecommons.org/licenses/by/4.0/
Original: https://www.maxresearchcollective.com/articles/nist-and-trustworthy-ai.html
Cite as: Newell, J. (2026). NIST and Trustworthy AI: The Lab That Never Fines Anyone. MAX Research Collective, ART-HLC-022. https://www.maxresearchcollective.com/articles/nist-and-trustworthy-ai.html

A lab with no power to fine

The National Institute of Standards and Technology (NIST) is the federal government's measurement laboratory, 125 years old this year and part of the Department of Commerce. It can't fine a company. A NIST document binds nobody until someone with power adopts it — a law, a regulator, a buyer or an accreditor — and most of its work has come to matter that way.

Its older standards protect systems that work as designed. Its AI Risk Management Framework (AI RMF), released in January 2023 and still at version 1.0, covers the harm those standards never had to: a model that's secure, patched and still wrong, or worse for one group of people than another. Laws have started pointing to it, unevenly. Texas names NIST's Generative AI Profile as a defense, Montana requires one class of AI deployer to consider the framework, and Colorado wrote in a defense for following it, then repealed that defense before it took effect.

Healthcare raises the stakes most. Hospitals and federal buyers already ask AI vendors for the evidence the framework describes, and accreditors put the same

questions to the hospitals. The same pressure reaches the apps and wearables people use at home, through federal and state rules that never name NIST and still ask for what it describes.

What NIST is

NIST is the United States' measurement laboratory. It began in 1901 as the National Bureau of Standards, so that a kilogram, a volt and an hour meant the same thing in every factory and courtroom. Congress renamed it in 1988 and widened its job to measurement science, standards and technology. It has two campuses, in Gaithersburg, Maryland, and Boulder, Colorado. Four of its staff have won the Nobel Prize in physics.

Its influence comes from how it works. NIST publishes a draft and takes public comment, then revises and publishes again; the result is free to anyone. Congress writes it into statutes and agencies cite it in rules. Buyers put it in contracts and auditors measure against it, long after NIST has moved on to the next draft.

Most people meet it without knowing. The clock on a phone traces back to its atomic clocks, and the padlock in a browser rests on encryption standards it published (Table 1).

Table 1. NIST documents that became binding somewhere, long before AI.

PUBLICATION	WHAT IT IS	WHERE IT ENDS UP BINDING
FIPS 197 and FIPS 203 to 205	AES encryption (2001) and the first post-quantum encryption and signature standards (2024)	Required for federal systems; built into nearly every browser, phone and bank
FIPS 140-3	Testing of cryptographic modules	Federal purchasing, and many health and finance contracts
SP 800-53 Rev. 5	The catalog of security and privacy controls	Federal systems under FISMA, and FedRAMP cloud services
CSF 2.0	The Cybersecurity Framework (1.0 in 2014, 2.0 in February 2024)	Used across industry; HHS has to take it into account when it penalizes a HIPAA breach
SP 800-66 Rev. 2	Implementing the HIPAA Security Rule (February 2024)	The practical guide covered entities use for HIPAA security
Privacy Framework 1.0	Managing privacy risk (2020)	Voluntary; widely used in privacy programs
SP 800-63-4	Digital identity, or proving who someone is online (2025)	Federal login systems; patient portals often follow it

A FIPS (Federal Information Processing Standard) is one that federal systems must meet; an SP (Special Publication) belongs to NIST's guidance series. FISMA, the Federal Information Security Modernization Act, applies SP 800-53 to federal systems, and FedRAMP authorizes cloud services for federal use. HIPAA is the Health Insurance Portability and Accountability Act, which HHS (the Department of Health and Human Services) enforces. Figure 1 draws the four routes those documents took from voluntary to required.

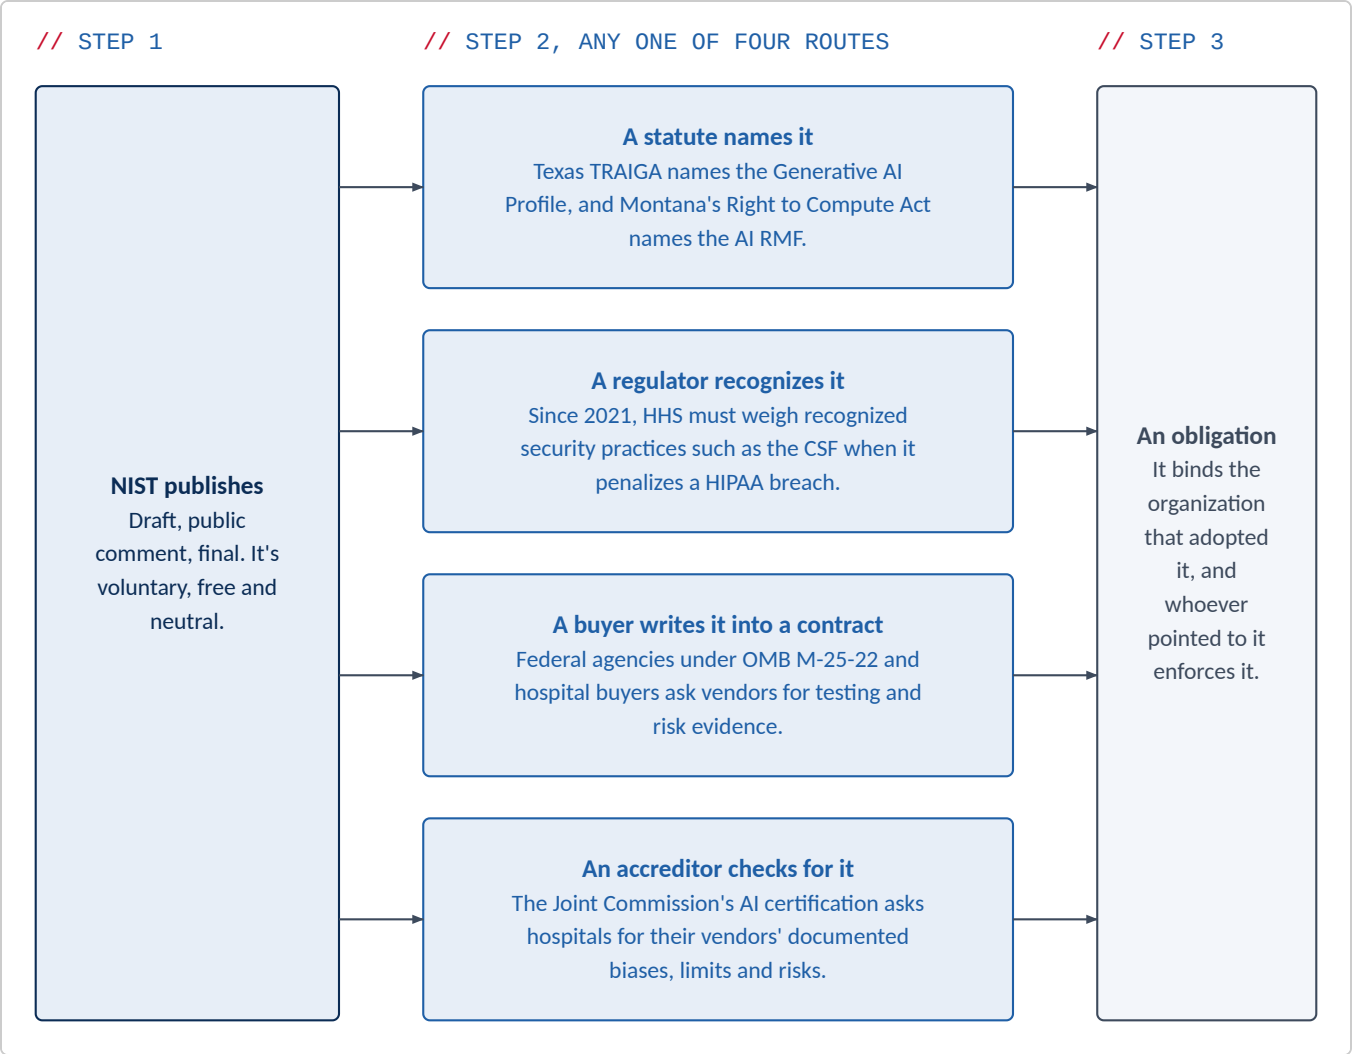


Figure 1. Every route from a voluntary NIST document to an obligation runs through someone else, and NIST enforces none of them.

What its AI work adds

Every document in Table 1 assumes the system works as designed and asks how to protect it, and AI breaks that assumption. A fully patched model can still answer with confidence and be wrong. It can work worse for one group than another, or drift as its data changes, and no encryption standard has anything to say about either.

The AI RMF, published as NIST AI 100-1, is written for that. It treats AI as socio-technical, meaning the risk depends on the model, the data, the people using it and how far they trust it. Like the Cybersecurity Framework, it asks for outcomes and names no products. What changed is the question (Table 2).

Table 2. What NIST's security and privacy frameworks ask, and what the AI RMF adds.

SECURITY AND PRIVACY FRAMEWORKS ASK	THE AI RMF ALSO ASKS
Can someone break in, steal it or alter it?	Is the output right, and how do we know?
Is personal data collected, kept and shared properly?	Does it work as well for everyone it's used on?
Can we detect an attack, respond and recover?	Can a person see why it said what it said, and overrule it?
	Will people trust it more than they should?
	What happens to people outside the organization when it fails?

Congress asked for the framework in the National AI Initiative Act, law since January 2021, and NIST released it two years and two public drafts later. Its text has stayed the same while the politics moved. America's AI Action Plan, published in July 2025, directed NIST to revise it and take out its references to misinformation, to diversity, equity and inclusion, and to climate change. NIST says the revision is under way; by 3 November 2026 nothing had been published, and the state laws that name the framework still cite its current title (Figure 2).

// EVERY DATE, NIST OR LAW

1901	●	National Bureau of Standards founded.
1988	●	Renamed the National Institute of Standards and Technology.
2014	●	Cybersecurity Framework 1.0: the outcome-based model the AI RMF later follows.
Jan 2021	□	National AI Initiative Act directs NIST to write a voluntary AI risk framework. The same month, Public Law 116-321 tells HHS to weigh recognized security practices in HIPAA enforcement.
Jan 2023	●	AI RMF 1.0 (NIST AI 100-1) and its Playbook released after two drafts and public workshops.
Feb 2024	●	CSF 2.0 adds a GOVERN function, the word the AI RMF put at its center.
Jul 2024	●	Generative AI Profile (NIST AI 600-1): twelve risks and about 210 suggested actions.
Nov 2024	●	NIST AI 100-4 on synthetic content: labeling, provenance and detection.
Mar 2025	●	NIST AI 100-2 E2025: the taxonomy of attacks on AI, prompt injection and data poisoning among them.
Apr 2025	□	OMB M-25-21 and M-25-22: minimum practices for high-impact federal AI, and what vendors must supply.
Jul 2025	□	America's AI Action Plan directs NIST to revise the AI RMF and remove its references to misinformation, diversity, equity and inclusion, and climate change.
Dec 2025	●	Cyber AI Profile (NIST IR 8596), preliminary draft: CSF 2.0 applied to securing AI.
Jan 2026	□	Texas TRAIGA takes effect and names the Generative AI Profile as a defense.
May 2026	□	Colorado SB 26-189 repeals the 2024 AI Act and its defense for following the AI RMF; the new law starts 1 January 2027.
Jul 2026	●	NIST AI 300-1 draft: a template for public-facing AI documentation, the model card.
Sep 2026	□	EO 14434 tells agencies to write "Super Intelligence" in place of "Artificial Intelligence" in new documents; in October NIST renames its AI control overlay project COSIS.
Oct 2026	□	ONC HTI-5 final rule takes the decision-support model card out of health IT certification.
Ahead	◌	AI RMF revision. NIST says it's under way, with no draft, version or date published.

● NIST a NIST publication or program

□ LAW a law, order, rule or policy that touches it

◌ NEXT announced, not yet published

Figure 2. The tag beside each date says whether NIST published it or a law, order, rule or policy acted on it; dates are publication or effective dates.

Four functions, nineteen categories

The core of the framework is four functions. GOVERN sets policy, names who's accountable and keeps control of suppliers, and it applies throughout a system's life. MAP establishes what the system is for and who it affects; MEASURE tests it before and after release; MANAGE decides what to fix, watches it in use and retires it safely. Those three repeat each time the system changes (Figure 3).

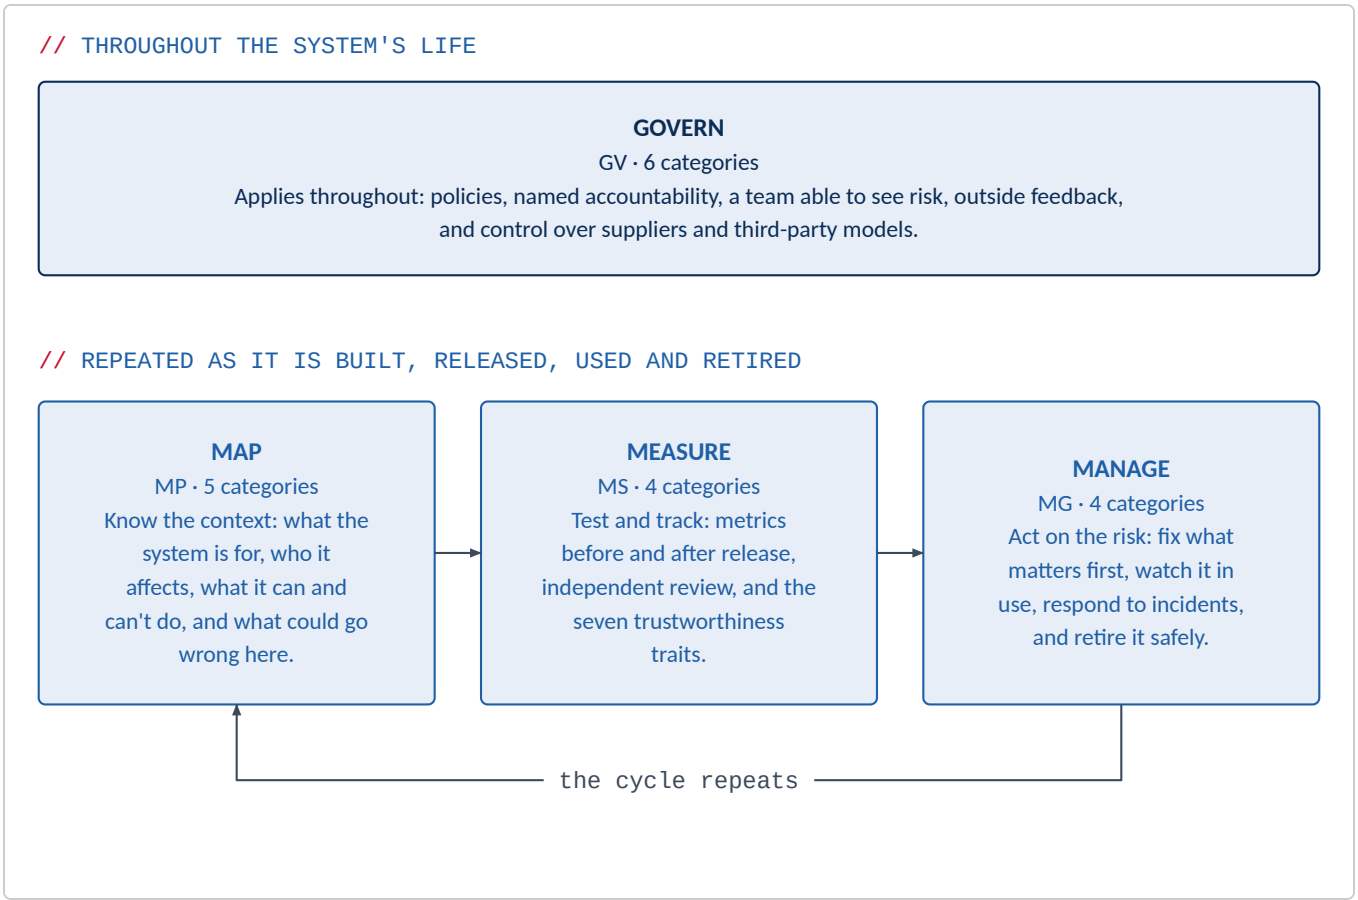


Figure 3. GOVERN sits over the whole life of the system, and the other three functions loop back to MAP each time the system changes.

The functions hold 19 categories and 72 outcomes, which NIST calls subcategories, spread unevenly. MEASURE 2, the evaluation against the trustworthiness traits, carries 13 outcomes on its own, as many as the whole MANAGE function (Figure 4). Appendix D lists every category.

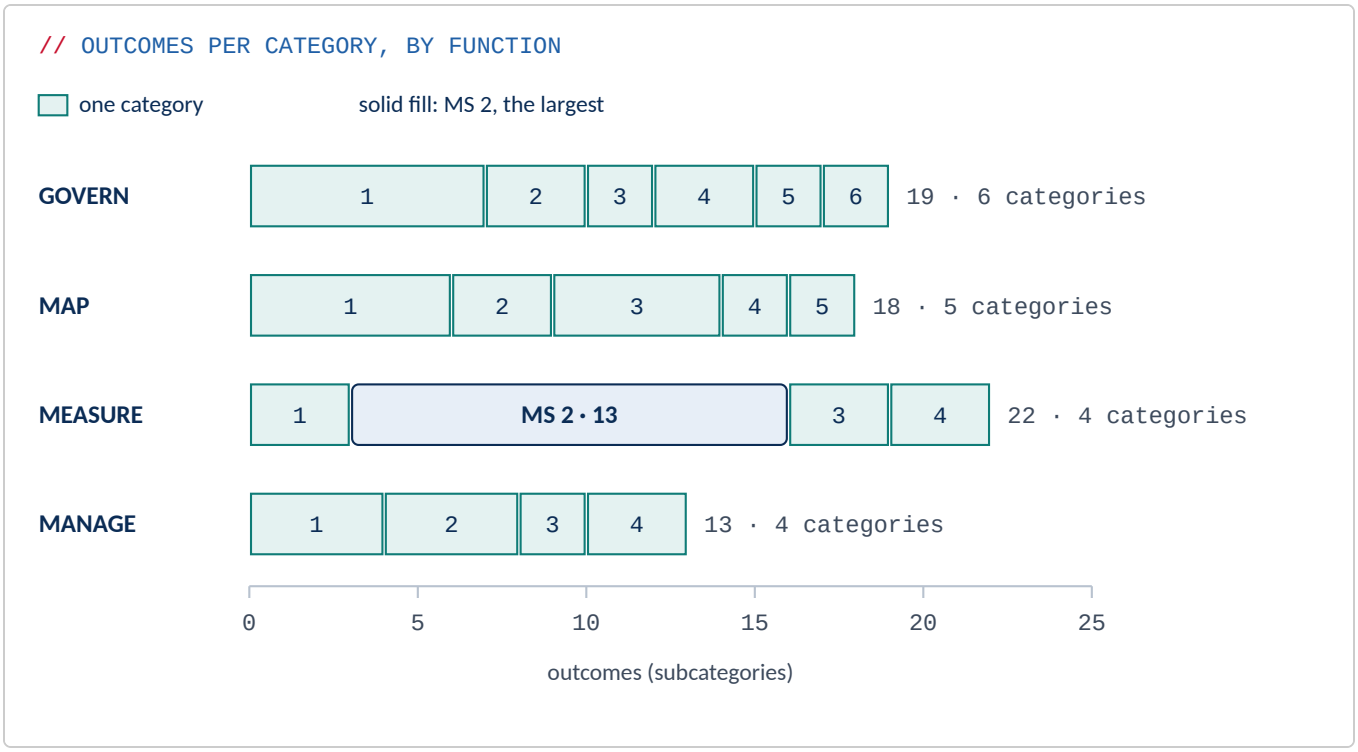


Figure 4. Each segment is one category, and the pale segment, MEASURE 2, holds more outcomes than any other.

Seven traits of trustworthy AI

The framework defines ”trustworthy” with seven traits (Figure 5). Valid and reliable is the base, since a system that gives wrong answers is unsafe whatever else it does well. Accountable and transparent spans the rest: safe, secure and resilient, explainable and interpretable, privacy-enhanced, and fair with harmful bias managed.

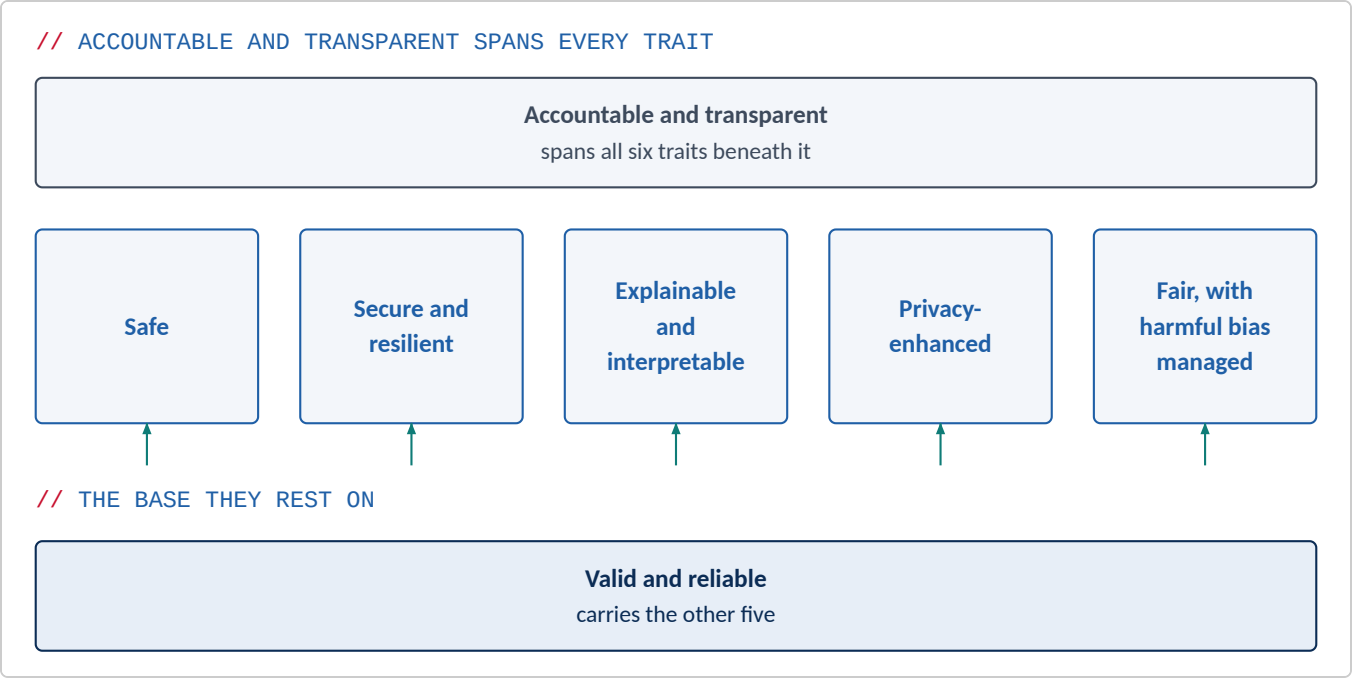


Figure 5. Valid and reliable carries the other five, and accountable and transparent spans every one of them.

The traits pull against each other. A more explainable model may be less accurate, and stronger privacy can leave less data to test for bias. The framework asks organizations to make those trade-offs on purpose and write them down, and most of them get made during training, by people who never recorded them.

The Generative AI Profile

A profile applies the framework to one kind of system. NIST AI 600-1, published in July 2024, does it for generative AI: chat assistants, image and voice generators, and the large models behind them. It names twelve risks and lists about 210 suggested actions mapped to 49 of the 72 subcategories. It's voluntary, and it's also the only NIST AI profile a state statute names: Texas's.

Confabulation is NIST's word for a fluent, confident answer that's false, which most people call a hallucination. Human-AI configuration covers the ways people and a system go wrong together: leaning on it too hard, deferring to it out of habit, or treating it as a person. The tags under each risk in Figure 6 mark where it bites hardest, and they're the Collective's reading of the profile; NIST assigns none.

// TWELVE RISKS, IN NIST'S ORDER

RISK 1
CBRN information
Easier access to chemical, biological, radiological or nuclear weapons knowledge.
Tags: Largest general models

RISK 2
Confabulation
Confident, fluent answers that are false: invented facts, citations or doses.
Tags: Consumers, Healthcare

RISK 3
Dangerous, violent or hateful content
Output that incites, threatens or encourages self-harm.
Tags: Consumers

RISK 4
Data privacy
Leaking, inferring or re-identifying personal data, health data included, from training data or prompts.
Tags: Consumers, Healthcare

RISK 5
Environmental impacts
The energy and water used to train and run large models.
Tags: Society at large

RISK 6
Harmful bias or homogenization
Worse results for some groups, and everyone nudged toward the same answers.
Tags: Consumers, Healthcare

RISK 7
Human-AI configuration
Over-reliance, automation bias, and treating the system as a person.
Tags: Consumers, Healthcare

RISK 8
Information integrity
False or misleading content made cheap to produce and hard to trace.
Tags: Consumers

RISK 9
Information security
Prompt injection, data poisoning, model theft, and help for attackers.
Tags: Consumers, Healthcare

RISK 10
Intellectual property
Reproducing copyrighted work, or someone's likeness or voice, without rights.
Tags: Consumers

RISK 11
Obscene, degrading or abusive content
Including sexual images of real people made without consent.
Tags: Consumers

RISK 12
Value chain and component integration
Third-party models, data and tools the deployer can't see into.
Tags: Healthcare

Figure 6. Nine of the twelve risks carry a Consumers tag and six a Healthcare tag, and five carry both.

The rest of the family

Several of the framework's companion documents are still drafts, and a final document can differ from its draft (Table 3).

Table 3. The NIST AI family, with each document's status on 3 November 2026.

DOCUMENT	WHAT IT ADDS	STATUS
AI 100-1	The AI RMF itself, with its Playbook of suggested actions	Final, 1.0; revision under way
AI 600-1	Generative AI Profile: twelve risks, about 210 actions	Final
AI 100-2 E2025	Adversarial machine learning: how AI is attacked, in shared terms	Final
AI 100-4	Synthetic content: labeling, provenance, detection	Final
SP 800-218A	Secure software development for generative AI and its training data	Final
SP 1270	Where bias in AI comes from: systemic, statistical and human	Final
IR 8596	Cyber AI Profile: CSF 2.0 outcomes for securing AI	Preliminary draft
AI 300-1	Public-facing AI documentation: a model card template	Initial draft
AI 800-1	Managing misuse risk of the most capable foundation models	Second draft
COSIS	SP 800-53 control overlays for assistants, predictive AI and agents	Outline stage

A model card is a short public document that says what a model is for, what it was tested on and where it fails. An overlay tailors the SP 800-53 controls to one kind of system.

The framework can't give anyone a certificate. Nobody can be "NIST AI certified", and a product that claims it is overstating. The certifiable standard is ISO/IEC 42001, a management system for AI that an accredited body audits, and NIST publishes a crosswalk matching each AI RMF outcome to the 42001 clause that covers it.

Why consumers live with it

Nobody using a chatbot or a symptom checker reads the AI RMF, and its effects reach them anyway. It's the closest thing the United States has to an agreed description of what a careful AI maker does: someone accountable for the system, testing on people like the ones using it, written limits, and watching after release with the power to turn it off. When a regulator or a court asks whether a company took reasonable care, this framework is more and more often the yardstick.

It gives an individual no right to sue. It protects people through whatever law or contract points to it, so the useful questions for a consumer are the framework's own:

1. Where does my data go? It matters whether it leaves the device, trains the model, and can be deleted.
2. How does it say when it's unsure? A careful system gives its confidence and its sources.
3. Was it tested on people like the ones using it? Age, language, accent, skin tone and health conditions all change how well AI works.
4. Who decides? For health, money or legal rights, a qualified person should make the call with the AI's evidence in front of them.
5. Can I reach a person, and appeal? A complaint route that leads nowhere is a failure of GOVERN.
6. What backs a claim of "NIST compliant"? NIST issues no AI certificate, so the follow-up asks what was tested and whether anyone outside checked it.

Healthcare and the other regulated industries

Regulated industries already answer to several regulators at once. AI leaves every one of them in place and adds risks none of them was written for, and the AI RMF is what they increasingly use to talk about those risks.

One AI tool in a clinic can fall under HIPAA as a handler of patient data. The Food and Drug Administration (FDA) may treat it as a medical device, and the Office of the National Coordinator for Health IT (ONC), now within HHS's Assistant Secretary for Technology Policy, may certify it as part of a health record system. State AI and medical practice laws, hospital accreditation and the buyer's contract come on top (Figure 7).

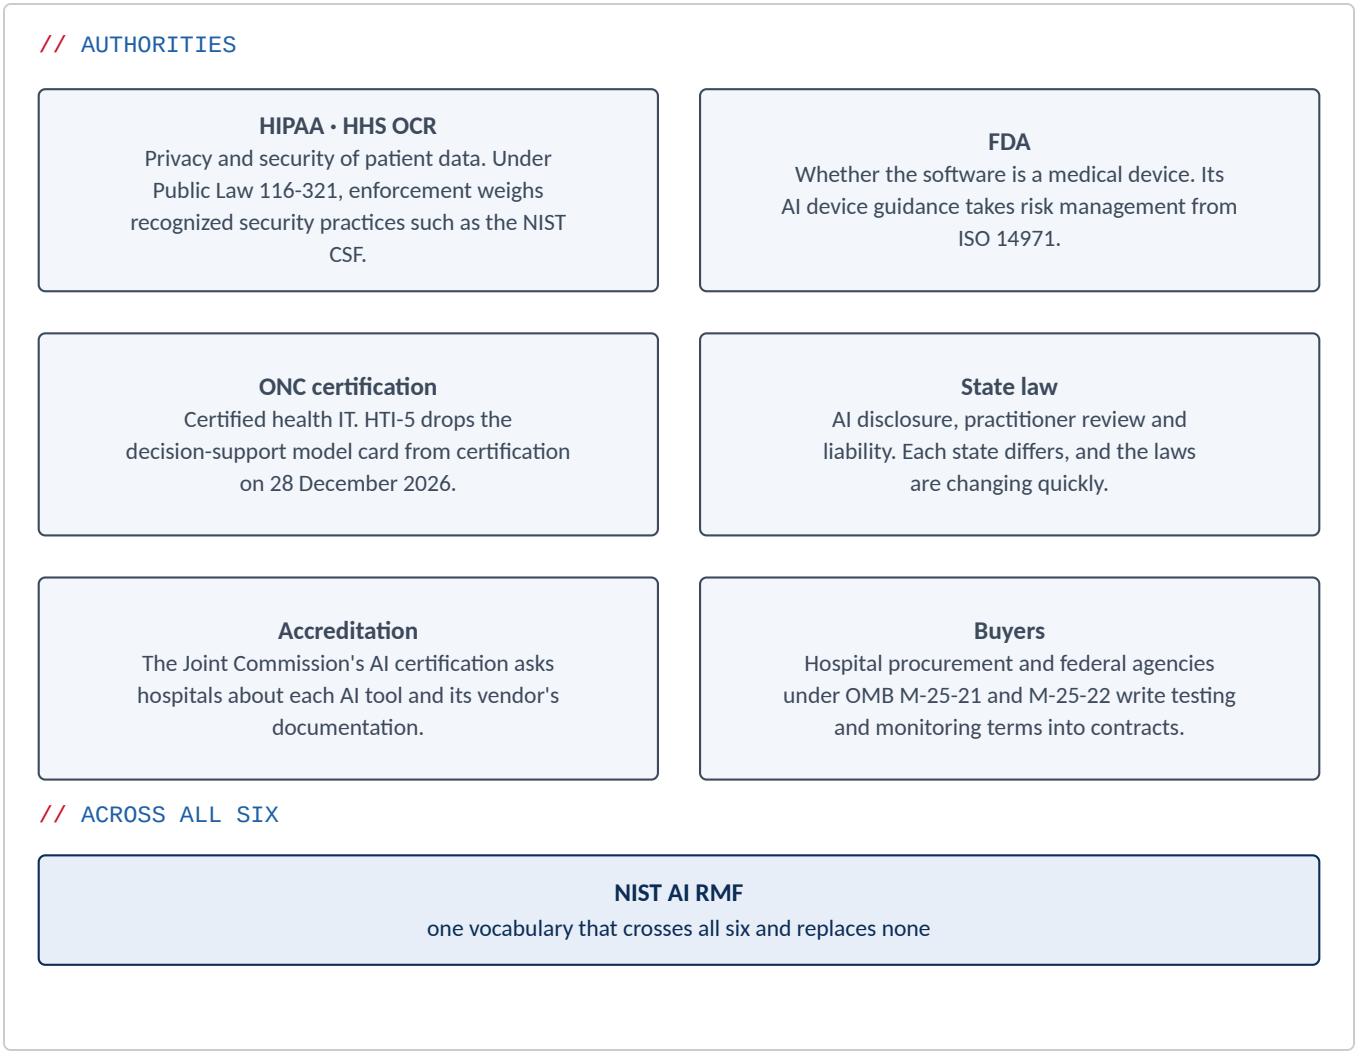


Figure 7. Six authorities can reach one clinical AI tool at once, and the AI RMF runs under all of them without replacing any.

Where the framework is pointed to today

Each row of Table 4 is an instrument that names the framework, relies on it, or chose another route. Every status is as it stood on 3 November 2026, rows marked secondary rest on commentary, and none of it is legal advice.

Table 4. Instruments that name the AI RMF, rely on it, or chose another standard.

INSTRUMENT	EFFECT OF FOLLOWING THE FRAMEWORK	DETAIL
Texas TRAIGA, HB 149, § 552.105(e)(2)(D); in force 1 January 2026	Liability defense	A defendant "may not be found liable" if it substantially complies with the most recent Generative AI Profile or another recognized framework. It's the only statute in force that names a NIST AI profile.
Montana Right to Compute Act, MCA 2-10-205	Must be considered	A deployer of AI that controls critical infrastructure must have a risk policy that considers the AI RMF, ISO/IEC 42001 or an equivalent.
Colorado AI Act, SB 24-205, replaced by SB 26-189 (signed May 2026); new law from 1 January 2027	Defense repealed	The 2024 act's defense for following the AI RMF never took effect. The replacement names no framework, exempts HIPAA covered entities and their business associates except for employment decisions, and still requires a general notice to patients. Enforcement is stayed in X.AI v. Weiser (secondary).
Connecticut, Public Act 26-15	Unverified	Commentary reports a safe harbor for following the AI RMF or ISO/IEC 42001; the act's text couldn't be read to confirm it (secondary).
Federal agencies as buyers, OMB M-25-21 and M-25-22	Evidence buyers ask for	Patient diagnosis, risk assessment and treatment are presumed high-impact, which brings pre-deployment testing, an impact assessment, monitoring, human oversight and appeals. The memos cite no NIST document, and the evidence they call for matches the AI RMF's MEASURE and MANAGE outcomes.

INSTRUMENT	EFFECT OF FOLLOWING THE FRAMEWORK	DETAIL
Joint Commission, RUAIH guidance 2025 and certification 2026	Evidence buyers ask for	The guidance, written with CHAI, names the AI RMF among its sources. The certification asks hospitals for "vendor documentation of known biases, limitations, and risks".
ONC HTI-5, 91 FR 67916; effective 28 December 2026	Requirement removed	Certification drops the decision-support model card attributes and the risk-management requirement. Hospitals can still ask for both in contracts.
FDA, AI-enabled device draft guidance, January 2025	Uses another standard	Device risk management follows ISO 14971. An AI RMF program supports it and doesn't stand in for it.
HIPAA enforcement, Public Law 116-321	Mitigation, for security	HHS must consider recognized security practices in place for the previous 12 months when it sets penalties. This covers the CSF, and it shows the route the AI RMF could follow.

TRAIGA is the Texas Responsible Artificial Intelligence Governance Act. OMB is the White House Office of Management and Budget, whose memos bind federal agencies, and HTI-5 is ONC's fifth Health Data, Technology, and Interoperability rule. CHAI, the Coalition for Health AI, wrote the Joint Commission's Responsible Use of AI in Healthcare guidance (RUAIH) with it. A covered entity is a provider, health plan or clearinghouse; a business associate handles patient data on its behalf.

Why the stakes run higher in healthcare

- The person harmed is someone other than the user. A clinician uses the tool, and the patient bears the error. MAP asks for both to be written down.
- Automation bias is a clinical hazard. Busy professionals come to follow a tool that's usually right, including when it's wrong. NIST lists that under human-AI

configuration, and the safe design leaves every medical decision to a trained person holding evidence they can audit.

- Bias shows up as worse care. A model that works less well for one population delivers worse care to it. MEASURE 2.11 asks for fairness to be measured, and healthcare has the data, provided small groups stay large enough that no patient can be re-identified.
- Patient data can't be taken back. A leaked diagnosis stays leaked. Running inference (the step where a trained model produces its answer) where the data already lives, on the device or inside the clinic, removes a class of risk that other designs have to manage.
- Evidence has to last. Clinical decisions get reviewed years later, and a tamper-evident record of what the tool showed and what the professional decided is the evidence that review will ask for.

Banks, insurers, utilities and federal contractors face the same pattern. Each already runs model-risk, security or safety programs under its own regulators, and the AI RMF gives those programs one AI vocabulary.

Health apps, wearables and AI chat

Most health software people touch never enters a hospital: a fitness tracker, a sleep or cycle app, a blood-pressure or glucose wearable, a symptom checker, an AI companion someone talks to at night. Which rules reach it depends on who offers it and what it claims to do. Those rules never cite the AI RMF, though each asks for something the framework already describes.

HIPAA is the first surprise for most users. It covers health plans, providers and the companies working for them, so a hospital's portal app is inside it. An app the patient picks from a store usually isn't. When a patient sends records from the portal to that app, HIPAA's protection stays behind with the portal (Table 5).

Table 5. Who reaches health and wellness software, as of 3 November 2026.

SOFTWARE	WHO REACHES IT	WHAT IT ASKS FOR
Portal or app offered by a provider or health plan	HIPAA, through HHS	Privacy and security of patient data, and breach notice
Consumer health app or connected device	FTC: the Health Breach Notification Rule as amended in 2024, and Section 5 of the FTC Act	Notice to users, the FTC and sometimes the media after a breach of identifiable health data; no deceptive privacy or accuracy claims
The same app, in some states	Washington's My Health My Data Act, in force since 2024, and similar laws in Nevada and Connecticut	Consent to collect consumer health data and separate consent to share it; Washington's law carries a private right of action
Wearable estimating blood pressure, glucose or heart rate variability	FDA, General Wellness guidance revised 6 January 2026	A non-invasive estimate can stay a wellness product if it names no disease, gives no clinical threshold and prompts no clinical action; an invasive sensor stays a device
Software that recommends to clinicians	FDA, clinical decision support guidance revised January 2026	The clinician must be able to review the basis of each recommendation independently, or the software is a device
AI therapy and companion chatbots	Illinois Wellness and Oversight for Psychological Resources Act (August 2025); California SB 243 and New York's companion-model law (in force); Tennessee SB 1580 (July 2026); Oregon SB 1546 (from 2027)	Illinois bars AI from therapeutic decision-making, with fines up to $10,000; the others require disclosure that the user is talking to AI, crisis protocols and protections for minors

Three of those rows moved in 2026. The FDA's January rewrite of its wellness and decision-support guidance let more wearables stay wellness products, which keeps them out of device review and puts the weight on what their makers claim. On 9 September 2026 the Federal Trade Commission (FTC) withdrew its 2021 policy statement on health apps. It called the statement obsolete, since the 2024 rule already covers health apps and connected devices like fitness trackers, and commentators read the withdrawal as narrowing what counts as a breach. The states keep moving on chatbots, and one bigger privacy law stalled: New York's legislature passed a broader health privacy act, and the governor vetoed it on 19 December 2025.

The framework matters to these products because each rule turns on evidence it organizes. The FDA's wellness line is an intended-use question, and that's MAP's first job. A wearable built to stay outside device territory needs its limits written as carefully as a device's. A companion chatbot's crisis protocol is MANAGE, incident response for a conversation that goes wrong. Confabulation in a symptom checker becomes an invented dose, and the FTC's deception authority reaches an app that claims accuracy it never measured, or a "NIST compliant" badge with no testing behind it.

What a working program keeps

The framework prescribes no tools, and organizations that use it well still end up keeping the same records (Table 6). Each answers a question a regulator, a buyer or a court will eventually ask.

Table 6. The records a working AI RMF program keeps.

RECORD	THE QUESTION IT ANSWERS	FRAMEWORK
AI inventory	Which AI systems do we run, from whom, and where?	GV 1.6
Model card for each system	What is it for, what was it tested on, and where does it fail?	MP 2.2, MS 2.9
Evaluation set	How accurate is it, and is it still that accurate?	MS 2.1, MS 2.5
Bias measurement	Does it work as well for every group it's used on?	MS 2.11
Adversarial test suite	What happens when someone tries to make it misbehave?	MS 2.7, AI 100-2
Risk register	What could go wrong, how likely, how bad, and who owns it?	MP 5.1, MG 1.2
Human oversight design	Who decides, with what evidence, and how do they overrule?	MP 3.5, GV 3.2
Supplier register	Which third-party models and data do we depend on, and on what terms?	GV 6.1, MG 3.1
Incident and feedback process	How do problems reach us, and what do we do about them?	GV 5.1, MG 4.3
Audit trail	Can we show, later, what the system showed and what a person decided?	GV 1.4, MS 2.8

An adversarial test suite is a set of deliberate attacks on the system, prompt injection among them, where hidden instructions in a document try to take over a model.

Table 7. What to watch after 3 November 2026.

WHEN	WHAT	WHY
28 December 2026	HTI-5 takes effect	The model card leaves health IT certification, and hospital contracts may take its place.
1 January 2027	Colorado SB 26-189 and Oregon SB 1546 start	A patient notice duty for covered entities, if the federal stay is lifted; companion-chatbot duties with statutory damages.
Not yet dated	The AI RMF revision	It may arrive under a new name after Executive Order 14434, which told agencies to write "Super Intelligence" in place of "Artificial Intelligence"; some categories may change.
Not yet dated	COSIS overlays	SP 800-53 controls for assistants and agents, likely to enter federal contracts.

Why it matters

Contracts, accreditation surveys and consumer-protection cases are making the AI RMF a working standard of care faster than any statute has. A vendor selling into a hospital will be asked for the framework's evidence whether or not a law requires it, because the Joint Commission asks the hospital and the hospital asks the vendor. An app maker selling straight to consumers meets the same questions from the FTC and the states, phrased as privacy, accuracy and crisis duties. An organization that already keeps the records in Table 6 hands them over, and the rest build them on a buyer's or a regulator's schedule.

The burden of proof moves too. A court, an attorney general or an accreditor can now read the framework and ask for the tests, the people they covered, the person who decided and the record that shows it. MARA, the Collective's Medical Agentic Resident Assistant, is being built to answer it that way, with the model on the clinic's own machine, a clinician making every decision and each step written to an audit log the Collective owns. For the patient on the other side of a clinical AI tool, or the person checking a wearable at two in the morning, the framework stays invisible,

and the part that matters is whether anyone can later show what the software knew and who acted on it.

Sources

1. NIST, AI Risk Management Framework, page read 3 November 2026.
    - nist.gov · ai risk management framework
2. NIST, AI 100-1, Artificial Intelligence Risk Management Framework (AI RMF 1.0), January 2023.
    - nvlpubs.nist.gov · NIST.AI.100 1
3. NIST, AI 600-1, Generative Artificial Intelligence Profile, July 2024.
    - nvlpubs.nist.gov · NIST.AI.600 1
4. NIST, Trustworthy and Responsible AI Resource Center, with the Playbook and the crosswalks.
    - airc.nist.gov
    - airc.nist.gov · crosswalks
5. The White House, America's AI Action Plan, July 2025.
    - whitehouse.gov · Americas AI Action Plan
6. The White House, Executive Order 14434, Inaugurating the Era of Super Intelligence, September 2026.
    - whitehouse.gov · inaugurating the era of super intelligence
7. NIST, IR 8596, Cyber AI Profile, preliminary draft, and the COSIS control overlays.
    - csrc.nist.gov · iprd
    - csrc.nist.gov · cosais
8. Texas Legislature, HB 149 (TRAIGA), enrolled.
    - capitol.texas.gov · html
9. Montana Code Annotated, 2-10-205.
    - mca.legmt.gov · section 0050

10. Colorado General Assembly, SB 26-189, and Holland & Knight, Colorado Governor Signs SB 189, May 2026.
    - leg.colorado.gov · bills
    - hklaw.com · colorado governor signs sb 189
11. OMB, M-25-21 and M-25-22.
    - whitehouse.gov · M 25 21 Accelerating Federal Use of AI through Innovation Governance…
    - whitehouse.gov · M 25 22 Driving Efficient Acquisition of Artificial Intelligence in…
12. ONC, HTI-5 final rule, 91 FR 67916, 28 October 2026.
    - govinfo.gov · pdf
13. The Joint Commission and CHAI, Responsible Use of AI in Healthcare, and the RUAIH certification guide.
    - digitalassets.jointcommission.org · dcfcf4f1a0cc45cdb526b3cb034c68c2
    - digitalassets.jointcommission.org · 7bd26620a3e54c0ead2138afa37fc167
14. FDA, Artificial Intelligence-Enabled Device Software Functions, draft guidance, January 2025.
    - fda.gov · download
15. FTC, FTC Withdraws Obsolete Policy Statement, 9 September 2026, and Epstein Becker Green, FTC Reverses Course on Health App Privacy, September 2026.
    - ftc.gov · ftc withdraws obsolete policy statement
    - ebglaw.com · ftc reverses course on health app privacy
16. Covington, FDA Issues Revised Guidance on General Wellness Products, January 2026, and King & Spalding, FDA Updates General Wellness and Clinical Decision Support Guidance.
    - cov.com · fda issues revised guidance on general wellness products
    - kslaw.com · fda updates general wellness and clinical decision support guidance…
17. Illinois Department of Financial and Professional Regulation, press release on HB 1806, 4 August 2025.

- idfpr.illinois.gov · 2025 08 04 idfpr press release hb1806

18. Orrick, 2026 State Chatbot Laws: Key Provisions and Regulatory Trends, April 2026.
    - orrick.com · 2026 State Chatbot Laws Key Provisions and Regulatory Trends
19. JD Supra, New York Governor Vetoes NY Health Privacy Act, December 2025, and the Future of Privacy Forum, health privacy comparison chart for New York, Washington and Connecticut.
    - jdsupra.com · new york governor vetoes ny health 4669037
    - fpf.org · FPF Health Comparison Chart NY WA CT
20. Secondary, for Connecticut Public Act 26-15 and X.AI v. Weiser: AI Compliance Atlas and AI Challenge Watch.
    - aicomplianceatlas.com · connecticut sb 2
    - aichallengewatch.com · xai v weiser

Laws, guidance and NIST's own documents are as they read on 3 November 2026. This is general information and not legal advice.

Tables 1 to 7 are complete data sets and take no appendix. Appendices A to G carry the data behind Figures 1 to 7, and Appendix H is the glossary.

Glossary

TERM	DEFINITION
AES (Advanced Encryption Standard)	The encryption method NIST published as FIPS 197 in 2001. It protects most stored and transmitted data in the world, which is why it's the usual first example of a NIST standard people use without knowing it.
AI RMF (AI Risk Management Framework)	NIST AI 100-1, the voluntary framework that sorts AI risk work into four functions, 19 categories and 72 outcomes. It's the document laws, buyers and accreditors now point to.
Automation bias	The habit of following a tool's output because it's usually right, including on the occasions it's wrong. In a clinic it turns a model's error into a clinician's decision.
Business associate	Under HIPAA, a company that handles patient data on behalf of a covered entity, such as a software vendor. It carries HIPAA duties of its own.
CDS (clinical decision support)	Software that gives clinicians recommendations about prevention, diagnosis or treatment. FDA guidance revised in January 2026 keeps it outside device rules when the clinician can review the basis of each recommendation independently.
CHAI (Coalition for Health AI)	A coalition of health systems, technology companies and others that writes guidance on using AI in care. It wrote the Joint Commission's 2025 responsible-use guidance with it.
Companion chatbot	An AI system built for ongoing personal conversation rather than a single task. California, New York, Oregon and others now require it to disclose that it's AI and to follow crisis protocols.
Confabulation	NIST's term for a fluent, confident answer that's false, often called a hallucination. It's risk 2 of the twelve in the Generative AI Profile.
COSIS	NIST's project to write SP 800-53 control overlays for AI assistants, predictive AI and agents, renamed in October 2026. Its overlays would tell federal agencies which security controls an AI system needs.

TERM	DEFINITION
Covered entity	Under HIPAA, a health care provider, health plan or clearinghouse. HIPAA's privacy and security rules apply to it directly.
Crosswalk	A table matching the parts of one framework to the parts of another. NIST's crosswalk to ISO/IEC 42001 lets one body of AI RMF work count toward both.
CSF (Cybersecurity Framework)	NIST's framework for managing security risk, version 1.0 in 2014 and 2.0 in February 2024. It's the outcome-based model the AI RMF follows, and HHS weighs it in HIPAA enforcement.
Deployer	The organization that puts an AI system to use, as distinct from the developer that built it. Several state laws place their duties on the deployer.
FedRAMP	The federal program that authorizes cloud services for use by agencies, built on SP 800-53. It's one of the routes by which a NIST control catalog becomes a condition of sale.
FIPS (Federal Information Processing Standard)	A NIST standard that federal systems are required to meet. FIPS 197 and FIPS 140-3 are the best known.
FISMA (Federal Information Security Modernization Act)	The law that requires federal agencies to secure their systems. It's what makes SP 800-53 binding on them.
Generative AI Profile	NIST AI 600-1, July 2024, which applies the AI RMF to generative AI and names twelve risks. Texas law names it as a defense.
FTC (Federal Trade Commission)	The federal agency that polices unfair and deceptive business practices under Section 5 of the FTC Act. It's the main federal regulator of health apps outside HIPAA.

TERM	DEFINITION
General wellness product	A low-risk product meant only to support a healthy lifestyle, which the FDA leaves outside device regulation. Its January 2026 guidance widened the category to some non-invasive estimates of blood pressure and glucose.
GOVERN, MAP, MEASURE, MANAGE	The four functions of the AI RMF core. GOVERN runs throughout, and the other three repeat across a system's life.
Health Breach Notification Rule	The FTC rule requiring vendors of personal health records, which since 2024 explicitly includes health apps and connected devices, to give notice after a breach of identifiable health data.
HIPAA (Health Insurance Portability and Accountability Act)	The federal law that governs the privacy and security of patient data. HHS's Office for Civil Rights enforces it.
HTI-5	The fifth of ONC's Health Data, Technology, and Interoperability rules, effective 28 December 2026. It removes the decision-support model card from health IT certification.
Human-AI configuration	Risk 7 of the Generative AI Profile, covering the ways people and an AI system go wrong together: over-reliance, automation bias, and treating the system as a person.
Inference	The step where a trained model produces an answer from new input. Where it runs decides where patient data has to travel.
ISO 14971	The international standard for risk management of medical devices. The FDA's AI device guidance uses it, so an AI RMF program supports it and doesn't replace it.
ISO/IEC 42001	The international standard for an AI management system, audited and certified by accredited bodies. It's the certifiable counterpart to the AI RMF.

TERM	DEFINITION
Model card	A short document stating what a model is for, what it was tested on and where it fails. NIST AI 300-1 drafts a template for one, and HTI-5 drops it from certification.
My Health My Data Act	Washington's consumer health privacy law, in force since 2024. It requires consent to collect health data outside HIPAA and separate consent to share it, and lets consumers sue.
OCR (Office for Civil Rights)	The part of HHS that enforces HIPAA. It's the office that weighs recognized security practices when it sets a penalty.
OMB (Office of Management and Budget)	The White House office whose memos set rules for federal agencies. M-25-21 and M-25-22 govern how agencies use and buy AI.
ONC (Office of the National Coordinator for Health IT)	The HHS office that certifies health IT, now part of the Assistant Secretary for Technology Policy. Its rules decide what certified electronic health records must do.
Overlay	A tailoring of the SP 800-53 control catalog for one kind of system. COSIS is writing overlays for AI.
Playbook	NIST's companion to the AI RMF, listing suggested actions for each outcome. It's guidance and carries no requirement of its own.
Profile	An application of the AI RMF to one kind of system or sector. The Generative AI Profile is the first, and a critical-infrastructure profile is planned.
Prompt injection	An attack that hides instructions in a document, page or message so that a model follows them in place of its user's. NIST AI 100-2 E2025 names it among the attacks on AI.
Risk register	The running list of what could go wrong with a system, how likely and how bad, and who owns each item. It answers MAP 5.1 and MANAGE 1.2.

TERM	DEFINITION
Safe harbor	A defense written into a law for anyone who meets a stated standard. Texas offers one for following the Generative AI Profile.
Socio-technical	NIST's description of AI risk as depending on the model, the data, the people using it and how far they trust it, together. It's why the framework asks about people as well as code.
SP (Special Publication)	NIST's main guidance series, such as SP 800-53 for security controls and SP 800-66 for the HIPAA Security Rule.
TRAIGA (Texas Responsible Artificial Intelligence Governance Act)	Texas HB 149, in force 1 January 2026. It's the only statute in force that names a NIST AI profile.

HEALTHCARE ART-HLC-022

// ABOUT THE AUTHOR

Jason Newell leads organizations and people in technology, and is a researcher, deep technologist and builder — from a CRAY to agentic AI. He founded MAX Research Collective to build AI for regulated industries that runs where its owner has control, with healthcare as its strongest focus. About Jason → Connect on LinkedIn →

Newell, J. (2026). MAX Research Collective, ART-HLC-022. CC BY 4.0