Securing AI Agent Systems: Recommendations in the NIST RFI Filings

Scott Wallsten and Sarah Oh Lam

June 2, 2026

2001 L Street NW ● Suite 500 ● Washington, DC 20036
Tel: (202) 828-4405 ● Email: info@techpolicyinstitute.org ● www.techpolicyinstitute.org

Securing AI Agent Systems: Recommendations in the NIST RFI Filings

Welcome to TPI's Docket Roundup, our semi-regular presentation of docket filings of interest to tech policy nerds. *DISCLAIMER: Filings are not affiliated with TPI. We do not necessarily agree with anything in these filings, but find them noteworthy to highlight.*

In this Issue: Over 80 substantive comments were filed in March 2026 with the National Institute of Standards and Technology (NIST) in reply to its Request for Information (RFI) on Securing AI Agent Systems (NIST-2025-0035). The RFI centered on novel, machine-learning-driven risks such as adversarial inference-time attacks, indirect prompt injection, data poisoning, model backdoors, and the potential for even uncompromised models to undermine confidentiality, integrity, or availability through specification gaming or misaligned objectives. NIST solicited input on whether these frameworks require agent-specific extensions related to threats, controls, assessment methods, deployment constraints, monitoring, research needs, and cross-sector collaboration. Docket: https://www.regulations.gov/search/comment?filter=NIST-2025-0035

Key Takeaways

1. Don't build a new regime — extend the existing one. This is the single most common position. Almost everyone (Amazon, Microsoft, ITI, CSET, Hitachi, USTelecom) wants NIST to add agent-specific *overlays/profiles* onto the AI RMF, SP 800-53, the Cybersecurity Framework, and SP 800-218A rather than create a parallel framework. Amazon's "security box" framing captures the mood: constraints as enabling infrastructure, not new bureaucracy.

2. Agents are treated as categorically different from both traditional software and chatbots. The recurring justification is the combination of non-deterministic reasoning + delegated authority + persistent memory + multi-step autonomous action at machine speed. BCG put it best: the security question shifts from *what a model can generate* to *what a system is permitted to execute*.

3. Indirect prompt injection is the consensus #1 threat — and model-level defenses are declared insufficient. Perplexity, the Frontier Model Forum, Elastic, and CrowdStrike all argue you need at least one *deterministic enforcement layer* outside the model. FAI/CSET cite adaptive attacks bypassing injection/jailbreak defenses at >90% success, so "the model will refuse" is explicitly rejected as a control.

4. Agent identity is its own sub-debate. A large cluster wants agents treated as first-class non-human identities with unique verifiable credentials, scoped authorization, and machine-speed revocation: Okta ("first-class digital identities"), Twilio ("agentic identity"), GoDaddy (domain-anchored Agent Name Service), OpenID Foundation, Cloudflare, and Ericsson (agents as "new insiders" under zero trust). Consensus that SP 800-63 is the foundation but incomplete for agents.

5. Least privilege + isolation/sandboxing + runtime governance is the most-recommended control stack. Booz Allen, Autodesk, Palo Alto, Red Hat, Splunk. Cognition AI (Devin/Windsurf) pushes specifically for VM-based isolation over containers because containers share the host kernel.

6. Continuous monitoring and observability are seen as essential but immature. Partnership on AI, Splunk, and CSET all flag that logging/observability *standards* are underdeveloped; Elastic proposes an "Agentic SOC" that baselines behavior, isolates suspicious agents, and rolls back before cascade.

7. Human oversight should scale with autonomy and stakes. Risk-tiered human-in-the-loop for consequential actions shows up in ServiceNow, Docusign (the "$1M renewal without a checkpoint" scenario), Intuit, and Okta.

8. Trade associations seek voluntary, risk-based, non-prescriptive guidance — backed by economic upside. CTA ($450–650B in annual revenue by 2030), CCIA, CTIA, USTelecom, BSA, TechNet, ITI, CEI, and ACT urge voluntary guidance. ACT (App Association) adds an important wrinkle: don't assume hyperscale capacity, or you'll lock out small developers.

9. Data provenance and integrity are reframed as security problems. Nielsen ("Data Nutrition Labels," provenance metadata), Cloudflare (cryptographic proof of permissions/provenance), OpenID (signed metadata), and HAI.AI's JACS (cryptographically signed agent communications, post-quantum via ML-DSA-87) frame data provenance as security matters.

10. Repeated calls for shared infrastructure. Agent-specific benchmarks (Google), a CVE-like AI vulnerability database and MITRE ATLAS extension (Lockheed Martin, AI Policy Network), reliability evaluation protocols (Princeton, arguing critical infra needs 99.9–99.999% vs. today's agents), updated incident definitions (Anthropic and Microsoft both say FISMA/SP 800-61 don't fit an *authorized* agent causing harm), and cross-sector info-sharing via ISACs (Agentic Futures Initiative) are proposals for shared infrastructure.

The risk frameworks cited across the filings included NIST AI RMF, SP 800-53, SP 800-218A (secure development), SP 800-207 (zero trust), MITRE ATLAS, OWASP Top 10 for Agentic Applications/LLMs, and ISO/IEC 42001. The most-repeated concrete asks were agent-specific profile/overlay; a deterministic enforcement/policy layer; agent identity + scoped credentials; least privilege across tools/APIs/code execution; runtime monitoring with rollback; risk-tiered human oversight; and standardized evaluation/assessment depth.

Will AI Security Rules Affect AI Economics?

A key question that is not explicitly part of the docket but nevertheless matters is how agent security standards will affect competition in AI. Compliance costs fall harder on smaller firms, and standards that codify one architecture advantage the firms whose products already match it. ACT (the App Association) is the only filer to raise this directly, warning that guidance built on hyperscale assumptions can lock out small developers.

That concern is worth taking seriously, but doesn't currently seem to be a major issue. For now, the market structure cuts against the capture story. The filings come from firms operating at many different layers of the agent stack, including model providers, identity vendors, cloud platforms, observability tools, and sandboxing infrastructure. No single layer is concentrated enough that one firm's preferences would define the framework. A fragmented ecosystem of this kind is the competitive condition that makes regulatory capture hard.

Instead, the standards process itself may pose a risk. Agent identity is the clearest example. A large cluster of filers wants agents treated as first-class identities with verifiable credentials and scoped authorization, and several are pushing their own approaches as the foundation. Whichever architecture NIST chooses becomes the default that every other firm has to build against. At that point, a technical standard could become a barrier to entry, and the firm whose design won the standard gains a durable advantage. The same dynamic applies to certification regimes and to any mandated enforcement layer.

Whether the framework is voluntary probably does not generally affect these points. Voluntary frameworks rarely stay voluntary in effect. A published NIST standard is likely to become the benchmark for reasonable diligence in negligence litigation and to be imported into federal procurement requirements and insurance underwriting criteria. Once that happens, conforming to the standard is no longer optional in any practical sense, and whatever competitive effects it carries are locked in with it. A standard can be voluntary in form and still shape the terms on which firms compete.

We are not claiming that agent security standards will necessarily help or harm competition. Instead, we are noting that the competitive and market-structure effects of these standards, which the docket largely sets aside in favor of threats and controls, deserve attention as the framework takes shape. As NIST develops its guidance, these economic effects are worth weighing alongside the security objectives. Whether and how standards affect entry, compliance costs, and market structure is a question the process should take up directly, even where the answers are not yet clear.

NIST RFI Overview

The NIST RFI asks stakeholders for concrete examples, best practices, case studies, and recommendations on securing AI agent systems that can take autonomous actions affecting external state. It frames agent security around novel ML-driven risks rather than ordinary software defects, naming adversarial inference-time attacks such as indirect prompt injection, data poisoning, model backdoors, and uncompromised models that still threaten confidentiality, integrity, or availability through specification gaming or misaligned objectives. NIST points respondents to existing foundations including NIST AI 100-2e2025, the AI RMF, the Generative AI Profile, NIST AI 800-1, SP 800-218A, and SP 800-53, while asking whether those frameworks need agent-specific extensions for threats, controls, assessment, deployment constraints, monitoring, research, and collaboration.

RFI Questions

1: Security threats, risks, and vulnerabilities affecting AI agent systems

2: Security practices for AI agent systems

3: Assessing the security of AI agent systems

4: Limiting, modifying, and monitoring deployment environments

5: Additional ecosystem, policy, research, and collaboration considerations

Table of Contents

Docket Links

Google

https://downloads.regulations.gov/NIST-2025-0035-0316/attachment_1.pdf
Questions addressed: primarily 1 and 2.

Amazon

https://downloads.regulations.gov/NIST-2025-0035-0532/attachment_1.pdf
Questions addressed: 1, 2, 3, 4, and 5.

Anthropic

https://downloads.regulations.gov/NIST-2025-0035-0460/attachment_1.pdf
Questions addressed: 1, 1(a), 2, 2(a), 2(c), 4(a), 4(c), 5, and 5(c).

OpenAI

https://downloads.regulations.gov/NIST-2025-0035-0504/attachment_1.pdf
Questions addressed: 1(a), 1(d), 1(e), 2(a), 2(e), 3(a), 3(b), 4(a), 4(b), 4(d), and 5(b)-(c).

Perplexity

https://downloads.regulations.gov/NIST-2025-0035-0505/attachment_1.pdf
Questions addressed: 1, 1(a), 1(b), 1(e), 2, 2(a), 3, 5, 5(a), and 5(c).

Microsoft

https://downloads.regulations.gov/NIST-2025-0035-0399/attachment_1.pdf
Questions addressed: 1, 1(a), 2, 3, 4, and 5.

Hugging Face

https://downloads.regulations.gov/NIST-2025-0035-0107/attachment_1.pdf
Questions addressed: 1, 1(a), 1(d), 1(e), 2, 2(a), 2(c), 2(e), 3, 3(a), 3(b), 3(c), 4, 4(a), 4(b), 4(d), 5, 5(a), 5(b), and 5(c).

Palantir

https://downloads.regulations.gov/NIST-2025-0035-0386/attachment_1.pdf
Questions addressed: 1, 1(a), 1(b), 1(e), 2, 2(a), 3, 4(b), 5, and 5(a).

Palantir Blog Attachment

https://downloads.regulations.gov/NIST-2025-0035-0386/attachment_2.pdf

Questions addressed: primarily 2, 4, and 5.

GitLab

https://downloads.regulations.gov/NIST-2025-0035-0133/attachment_1.pdf
Questions addressed: primarily 1, 2, 3, 4, and 5.

Cloudflare

https://downloads.regulations.gov/NIST-2025-0035-0344/attachment_1.pdf
Questions addressed: primarily 1, 2, 4, and 5.

JPMorgan Chase

https://downloads.regulations.gov/NIST-2025-0035-0345/attachment_1.pdf
Questions addressed: primarily 1, 2, 3, 4, and 5.

Salesforce

https://downloads.regulations.gov/NIST-2025-0035-0409/attachment_1.pdf
Questions addressed: primarily 1, 2, 4, and 5.

Intel

https://downloads.regulations.gov/NIST-2025-0035-0397/attachment_1.pdf
Questions addressed: 5.

Cisco

https://downloads.regulations.gov/NIST-2025-0035-0402/attachment_1.pdf
Questions addressed: 1, 2, 3, 4, and 5.

Ericsson

https://downloads.regulations.gov/NIST-2025-0035-0389/attachment_1.pdf
Questions addressed: 1, 2, 3, 4, and 5.

CrowdStrike

https://downloads.regulations.gov/NIST-2025-0035-0366/attachment_1.pdf
Questions addressed: primarily 1, 2, 4, and 5.

CrowdStrike

https://downloads.regulations.gov/NIST-2025-0035-0391/attachment_1.pdf
Questions addressed: 1(a), 1(d), 1(e), 2(a), 2(e), 3(a), 3(b), 4(a), 4(b), 4(d), 5(b), and 5(c).

HAI.AI / JACS Project

https://downloads.regulations.gov/NIST-2025-0035-0284/attachment_1.pdf

Questions addressed: 1, 2, 3, and 4.

Cognition AI

https://downloads.regulations.gov/NIST-2025-0035-0510/attachment_1.pdf

Questions addressed: 2(a), 2(b), 2(c), 2(d), 2(e), 3(a), 4(a), 4(b), and 4(d).

Agentic Futures Initiative

https://downloads.regulations.gov/NIST-2025-0035-0346/attachment_1.pdf

Questions addressed: 1, 1(a), 1(d), 2(a), 2(e), 3(a), 3(b), 4(a), 4(b), 4(d), 5, and 5(b).

Frontier Model Forum

https://downloads.regulations.gov/NIST-2025-0035-0349/attachment_1.pdf

Questions addressed: primarily 1, 3, and 5.

Partnership on AI

https://downloads.regulations.gov/NIST-2025-0035-0462/attachment_1.pdf

Questions addressed: 1(a), 1(b), 2(a), 2(c), 2(e), 3(b), 5(a), and 5(c).

OpenPolicy Coalition

https://downloads.regulations.gov/NIST-2025-0035-0340/attachment_1.pdf

Questions addressed: 1, 1(a), 1(d), 1(e), 2, 2(a), 2(e), 3, 3(a), 3(b), 4, 4(a), 4(b), 4(d), and 5.

OpenID Foundation

https://downloads.regulations.gov/NIST-2025-0035-0231/attachment_1.pdf

Questions addressed: 1, 2, 3, 4, and 5.

AI Policy Network

https://downloads.regulations.gov/NIST-2025-0035-0530/attachment_1.pdf

Questions addressed: 1, 1(a), 1(d), 2, 2(e), 3, 3(a), 3(b), 3(c), 4, 4(d), 5, 5(a), 5(b), and 5(c).

CTIA

https://downloads.regulations.gov/NIST-2025-0035-0392/attachment_1.pdf

Questions addressed: primarily 2(e), 4(c), and 5.

USTelecom

https://downloads.regulations.gov/NIST-2025-0035-0300/attachment_1.pdf

Questions addressed: primarily 1, 2, and 5.

Consumer Technology Association (CTA)

https://downloads.regulations.gov/NIST-2025-0035-0012/attachment_1.pdf

Questions addressed: primarily 1, 2, and 5.

The Digital Chamber

https://downloads.regulations.gov/NIST-2025-0035-0393/attachment_1.pdf

Questions addressed: 1, 1(a), 1(d), 2, 2(a), 2(e), 3, 3(a), 3(b), 4, 4(a), 4(b), 4(d), and 5.

BSA | The Software Alliance

https://downloads.regulations.gov/NIST-2025-0035-0457/attachment_1.pdf

Questions addressed: 1(a), 1(b), 2(a), 2(e), 5(a), 5(b), and 5(c).

Competitive Enterprise Institute (CEI)

https://downloads.regulations.gov/NIST-2025-0035-0473/attachment_1.pdf

Questions addressed: 4(b) and 5.

American Bankers Association and BPI/BITS

https://downloads.regulations.gov/NIST-2025-0035-0479/attachment_1.pdf

Questions addressed: primarily 1, 2, 4, and 5.

Computer & Communications Industry Association (CCIA)

https://downloads.regulations.gov/NIST-2025-0035-0320/attachment_1.pdf

Questions addressed: primarily 1, 2, 3, 4, and 5.

Software & Information Industry Association (SIIA)

https://downloads.regulations.gov/NIST-2025-0035-0367/attachment_1.pdf

Questions addressed: 1, 1(c), 2, and 5.

Foundation for American Innovation (FAI)

https://downloads.regulations.gov/NIST-2025-0035-0525/attachment_1.pdf

Questions addressed: 1(c), 2(a), 2(b), 3(a), 3(b), 4(b), 4(d), 5(a), and 5(b).

Center for Security and Emerging Technology (CSET)

https://downloads.regulations.gov/NIST-2025-0035-0335/attachment_1.pdf

Questions addressed: 2, 3, 4, and 5.

ACT | The App Association

https://downloads.regulations.gov/NIST-2025-0035-0410/attachment_1.pdf

Questions addressed: 1(a), 1(d), 2, 2(a), 2(e), and 3(b).

Information Technology Industry Council (ITI)

https://downloads.regulations.gov/NIST-2025-0035-0439/attachment_1.pdf

Questions addressed: 1, 2, 3, 3(b), 4, and 5.

TechNet

https://downloads.regulations.gov/NIST-2025-0035-0441/attachment_1.pdf

Questions addressed: primarily 1, 2, 3, 4, and 5.

Lockheed Martin

https://downloads.regulations.gov/NIST-2025-0035-0322/attachment_1.pdf

Questions addressed: 1, 1(a), 1(b), 2, 2(a), 2(b), 2(c), 2(e), 3, 3(a), 3(b), 3(c), 4, 4(a), 4(b), 4(e), 5, 5(a), and 5(b).

Palo Alto Networks

https://downloads.regulations.gov/NIST-2025-0035-0341/attachment_1.pdf

Questions addressed: 1, 2, 3, and 4.

Docusign

https://downloads.regulations.gov/NIST-2025-0035-0279/attachment_1.pdf

Questions addressed: 1, 1(a), 1(d), 2, 2(a), 2(e), 3, 3(a), 3(b), 4, 4(a), 4(b), and 4(d).

ServiceNow

https://downloads.regulations.gov/NIST-2025-0035-0234/attachment_1.pdf

Questions addressed: primarily 1, 2, 4, and 5.

Okta

https://downloads.regulations.gov/NIST-2025-0035-0260/attachment_1.pdf

Questions addressed: 1, 2, 3, and 4.

Twilio

https://downloads.regulations.gov/NIST-2025-0035-0357/attachment_1.pdf

Questions addressed: all RFI questions, 1(a)-(e), 2(a)-(e), 3(a)-(d), 4(a)-(e), and 5(a)-(e).

Wharton Accountable AI Lab

https://downloads.regulations.gov/NIST-2025-0035-0372/attachment_1.pdf

Questions addressed: 1(a), 1(d), 1(e), and 5(a).

Johns Hopkins APL / Institute for Assured Autonomy

https://downloads.regulations.gov/NIST-2025-0035-0381/attachment_1.pdf

Questions addressed: 1, 1(a), 1(e), 2, 2(a), 2(e), 3, 3(b), 4, 4(a), 4(b), and 4(d).

Princeton CITP / Princeton AI Lab Researchers

https://downloads.regulations.gov/NIST-2025-0035-0487/attachment_1.pdf

Questions addressed: 1, 1(d), 1(e), 2, 2(a), 2(e), 3, 3(a), 3(b), 4, 4(a), 4(b), 5, 5(b), 5(c), and 5(e).

2025 AI Agent Index Comment MIT

https://downloads.regulations.gov/NIST-2025-0035-0488/attachment_1.pdf

Questions addressed: 1(a), 1(d), 2(a), 3(a), 3(b), 4(a), 4(d), and 5(b).

The 2025 AI Agent Index Paper MIT

https://downloads.regulations.gov/NIST-2025-0035-0488/attachment_2.pdf

Questions addressed: primarily 2, 3, 4, and 5.

California Privacy Protection Agency

https://downloads.regulations.gov/NIST-2025-0035-0385/attachment_1.pdf

Questions addressed: primarily 3 and 5.

Leidos

https://downloads.regulations.gov/NIST-2025-0035-0364/attachment_1.pdf

Questions addressed: 1, 1(a), 2, 2(a), 2(e), 3, 3(a), 4, 4(a), 4(d), 5, 5(b), and 5(e).

Intuit

https://downloads.regulations.gov/NIST-2025-0035-0400/attachment_1.pdf

Questions addressed: primarily 1, 2, 3, 4, and 5.

Red Hat

https://downloads.regulations.gov/NIST-2025-0035-0450/attachment_1.pdf

Questions addressed: 1, 2, 3, 4, and 5.

Autodesk

https://downloads.regulations.gov/NIST-2025-0035-0459/attachment_1.pdf

Questions addressed: all RFI questions, 1(a)-(e), 2(a)-(e), 3(a)-(d), 4(a)-(e), and 5(a)-(e).

GoDaddy

https://downloads.regulations.gov/NIST-2025-0035-0486/attachment_1.pdf

Questions addressed: 1(a), 1(d), 1(e), 2(a), 2(d), 2(e), 3(a), 3(b), 4(c), 4(d), 4(e), 5(a), and 5(b).

Nielsen

https://downloads.regulations.gov/NIST-2025-0035-0502/attachment_1.pdf

Questions addressed: 1, 1(d), 2, 3, and 4.

Databricks

https://downloads.regulations.gov/NIST-2025-0035-0495/attachment_1.pdf

Questions addressed: 1(a)-(e), 2(a)-(e), 3(a)-(d), 4(a)-(d), and 5(a)-(c).

Elasticsearch Federal, Inc.

https://downloads.regulations.gov/NIST-2025-0035-0181/attachment_1.pdf

Questions addressed: 1, 2, 3, 4, and 5.

ISO/IEC JTC 1/SC 42 Japan National Body

https://downloads.regulations.gov/NIST-2025-0035-0216/attachment_1.pdf

Questions addressed: primarily 3.

Hitachi

https://downloads.regulations.gov/NIST-2025-0035-0313/attachment_1.pdf

Questions addressed: 1, 1(a), 1(d), 2, 2(a), 2(e), 3, 3(a), 3(b), 4, 4(a), 4(b), 4(d), and 5.

Consensys

https://downloads.regulations.gov/NIST-2025-0035-0411/attachment_1.pdf

Questions addressed: 1, 2, 3, 4, and 5.

Boston Consulting Group (BCG)

https://downloads.regulations.gov/NIST-2025-0035-0499/attachment_1.pdf

Questions addressed: primarily 1, 2, 4, and 5.

Splunk

https://downloads.regulations.gov/NIST-2025-0035-0535/attachment_1.pdf

Questions addressed: 1, 1(a), 1(d), 2, 2(a), 2(e), 3, 3(a), 3(b), 3(c), 3(d), 4(a), 4(b), 4(c), 4(d), 4(e), and 5.

Booz Allen Hamilton

https://downloads.regulations.gov/NIST-2025-0035-0374/attachment_1.pdf

Questions addressed: 1, 1(a), 1(b), 1(c), 1(d), 1(e), 2, 2(a), 2(b), 2(d), 2(e), 3, 3(a), 4, 4(a), 4(d), and 4(e).

PwC

https://downloads.regulations.gov/NIST-2025-0035-0351/attachment_1.pdf

Questions addressed: all RFI questions, 1(a)-(e), 2(a)-(e), 3(a)-(d), 4(a)-(e), and 5(a)-(e).

Lasso Security

https://downloads.regulations.gov/NIST-2025-0035-0458/attachment_1.pdf

Questions addressed: 1(a)-(e), 2(a), 2(b), 2(d), 2(e), 3(a)-(d), 4(a)-(e), 5, and 5(a).

18 Principles of Trust Framework

https://downloads.regulations.gov/NIST-2025-0035-0128/attachment_1.pdf

Questions addressed: primarily 1.

Aegis Protocol

https://downloads.regulations.gov/NIST-2025-0035-0507/attachment_1.pdf

Questions addressed: primarily 2 and 5.

10a Labs

https://downloads.regulations.gov/NIST-2025-0035-0388/attachment_1.pdf

Questions addressed: 1(a), 1(d), and 2(a).

10a Labs Prompt Intel Attachment

https://downloads.regulations.gov/NIST-2025-0035-0388/attachment_2.pdf

Questions addressed: primarily 1.

Identity Governance Consortium (IGC)

https://downloads.regulations.gov/NIST-2025-0035-0013/attachment_1.pdf
Questions addressed: 1(a), 1(c), 2(e), 3(a), 3(c), 4(c), 5(a), 5(b), 5(c), and 5(d).

Technological Society of Applied Research (TSAR)

https://downloads.regulations.gov/NIST-2025-0035-0057/attachment_1.pdf
Questions addressed: primarily 1, 2, 3, 4, and 5.

Foundation for Defense of Democracies (FDD)

https://downloads.regulations.gov/NIST-2025-0035-0342/attachment_1.pdf
Questions addressed: not specified in filing.

Frameworks Cited

NIST AI RMF

https://www.nist.gov/itl/ai-risk-management-framework
NIST AI RMF Trustworthy AI in Critical Infrastructure, April 7, 2026

NIST SP 800-63

https://pages.nist.gov/800-63-4/
NIST Digital Identity Guidelines, Revision 4, July 2025

NIST SP 800-218A

https://csrc.nist.gov/pubs/sp/800/218/a/final
NIST Secure Software Development Practices for Generative AI and Dual-Use Foundation Models, July 2024

NIST CSF 2.0

https://www.nist.gov/cyberframework
NIST NIST Cybersecurity Framework (CSF) 2.0, Feb. 2024

NIST SP 800-61

https://csrc.nist.gov/pubs/sp/800/61/r3/final
NIST Incident Response Recommendations and Considerations for Cybersecurity Risk Management, Revision 3, April 2025

NIST SP 800-207

https://csrc.nist.gov/pubs/sp/800/207/final

NIST Zero Trust Architecture, Aug. 2020

NIST SP 800-53

https://csrc.nist.gov/pubs/sp/800/53/r5/upd1/final

NIST Security and Privacy Controls for Information Systems and Organizations, Revision 5, Sept. 2020

NIST GenAI Profile

https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence

NIST Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, July 2024

NIST AI 100-2e2025

https://csrc.nist.gov/pubs/ai/100/2/e2025/final

NIST Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations, March 2025

NIST AI 800-1

https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.800-1.ipd.pdf

NIST Managing Misuse Risk for Dual-Use Foundation Models, U.S. AI Safety Institute, July 2024

OWASP Top 10 for Agentic Applications

https://genai.owasp.org/resource/owasp-top-10-for-agentic-applications-for-2026/

OWASP Top 10 for Agentic Applications for 2026, Dec. 2025

ISO/IEC 42001

https://blog.ansi.org/anab/iso-iec-42001-ai-management-systems/

ISO/IEC 42001: Artificial Intelligence Management Systems (AIMS), Aug. 2024

MITRE ATLAS

https://atlas.mitre.org/

MITRE ATLAS (Adversarial Threat Landscape for Artificial-Intelligence Systems)

HAI.AI/JACS

https://github.com/humanassisted/jacs

JSON Agent Communication Standard, generate and manage a persistent cryptographic identity for an agent. Post-quantum ready (pq2025 / ML-DSA-87) by default

NIST ML-DSA-87

https://nvlpubs.nist.gov/nistpubs/fips/nist.fips.204.pdf

NIST FIPS 204, Module-Lattice-Based Digital Signature Standard, Aug. 2024