NIST AI Risk Management Framework

Staffinity Alignment

NIST AI RMF 1.0 — January 2023

Prepared May 2026

security@staffinity.io | trust.staffinity.io | staffinity.io

EXECUTIVE SUMMARY

Staffinity aligns its AI governance practices with the NIST AI Risk Management Framework (AI RMF 1.0), published January 2023 by the National Institute of Standards and Technology. This document maps Staffinity's existing controls and processes to each of the four AI RMF core functions: **GOVERN, MAP, MEASURE,** and **MANAGE**. Staffinity uses this alignment to guide responsible AI deployment across client engagements and to demonstrate trustworthy AI practices to enterprise clients.

The NIST AI RMF is a voluntary, non-prescriptive framework that provides organizations with approaches to manage risks associated with AI systems. It is designed to be flexible and applicable across industries, AI technologies, and organizational contexts. Staffinity has adopted the AI RMF as a foundational governance reference and continuously improves alignment as the framework and accompanying playbooks evolve.

> **Note: The NIST AI RMF is a voluntary framework. Staffinity's alignment represents a commitment to responsible AI practices, not a formal certification. This document is accurate as of May 2026 and is subject to annual review.**

Document Scope

- AI System: Staffinity AI Agent Platform (workplace assistant deployed via Microsoft Teams)
- Model Provider: Anthropic Claude Sonnet (via customer API key)
- Infrastructure: AWS (ECS Fargate, RDS pgvector, ElastiCache Redis)
- Responsible Party: Brent Sudeck, CEO & AI Privacy and Security Officer (CISSP, CCSP, AWS SAA)
- Review Cycle: Annual (next review: May 2027)

AI RMF Core Functions at a Glance

Function	Purpose	Pages
GOVERN (GV)	Establish AI risk governance policies, accountability, and culture	3
MAP (MP)	Categorize and contextualize AI risk for each deployment	4
MEASURE (MS)	Analyze and assess AI system performance, trustworthiness, and risk	5
MANAGE (MG)	Prioritize and address AI risks; improve systems over time	6

GV | GOVERN

Establishing AI Risk Governance

Policies, accountability structures, and culture for managing AI risk across the organization. GOVERN encompasses the organizational structures, policies, and processes that establish accountability and oversight for AI systems throughout their lifecycle.

Staffinity Controls Mapping

GV-1.1	AI governance policy documented — Staffinity AI Governance Policy v1.0 (May 2026) defines scope, roles, risk tolerance, and review schedule.
GV-1.2	Brent Sudeck designated as AI Privacy & Security Officer with explicit authority over AI governance, risk management, and incident response.
GV-1.3	Roles and responsibilities documented across CEO, engineering, and AI operations; RACI matrix maintained in internal governance documentation.
GV-2.1	Organizational risk tolerance defined — client data protection is non-negotiable; acceptable risk includes model hallucinations within human-reviewed workflows.
GV-4.1	AI risk management integrated into deployment process: Yeager test suite (7-point), deployment checklist, and per-client sign-off required before go-live.
GV-5.1	Per-client deployment review conducted before go-live; contractual Data Processing Agreement (DPA) governs all data handling obligations.
GV-6.1	Governance policies reviewed annually; AI Governance Policy v1.0 includes explicit review schedule and change management process.
GV-7.1	Third-party AI provider (Anthropic) evaluated against NIST AI RMF criteria; Enterprise DPA enforces no-model-training clause; sub-processor list published at trust.staffinity.io/dpa.html.

MP | MAP

Categorizing and Contextualizing AI Risk

Understanding AI system context, intended use, and potential impacts before deployment. MAP activities establish awareness of the context surrounding each AI system and help identify risks before they materialize.

Staffinity Controls Mapping

MP-1.1	Each client deployment scoped to specific use cases documented in Exhibit A (Development Plan). Use cases outside scope require a formal change order and re-evaluation.
MP-2.1	Stakeholders identified per deployment: client employees (end users), client IT administrators, Staffinity (operator/processor), and Anthropic (model provider/sub-processor).
MP-2.3	AI system context documented — Staffinity agents operate as workplace assistants, not autonomous decision-makers. Human-in-the-loop is required for all consequential actions.
MP-3.1	Potential negative impacts assessed per deployment: hallucination risk (mitigated by RAG grounding on client documents), data leakage (mitigated by per-client AWS account isolation and PII detection), unauthorized access (mitigated by Microsoft Entra ID authentication + MFA).
MP-4.1	Benefits documented per client use case in Exhibit A: time savings on repetitive queries, error reduction through grounded responses, 24/7 availability for policy and procedure lookups.
MP-5.1	Deployment categorization: Staffinity agents are Level 3 (semi-autonomous) per the BVDW autonomy framework — routine tasks are automated; exceptions and high-risk queries are escalated to human supervisors.

MS | MEASURE

Analyzing and Assessing AI Risk

Ongoing evaluation of AI system performance, trustworthiness, and risk. MEASURE activities establish how AI system trustworthiness is assessed and tracked over time through metrics, testing, and monitoring.

Staffinity Controls Mapping

MS-1.1	Evaluation metrics defined pre-deployment: response latency (p95), error rate, circuit breaker trigger rate, tool call success rate. Baselines established during staging.
MS-2.1	Pre-deployment testing performed using Yeager 7-point test suite covering: authentication, tool execution, memory retrieval, response quality, PII detection, circuit breaker behavior, and multi-turn coherence.
MS-2.5	Bias and fairness assessment: agent system prompts reviewed for neutrality before deployment; agents instructed to avoid demographically differentiated responses; periodic spot-checks conducted.
MS-2.6	Explainability: agent responses are natural language with tool call transparency — SharePoint documents referenced by name, search results cited, reasoning visible to users in Teams.
MS-3.1	Performance monitored continuously via CloudWatch Synthetics canaries (1-minute interval), health check dashboards, CloudWatch Alarms with SNS notification, and ECS service metrics.
MS-4.1	Incidents logged in immutable audit trail (DynamoDB + S3 WORM Object Lock COMPLIANCE, 7-year retention). Circuit breaker triggers escalation email and internal HAL notification.
MS-4.2	Teams native thumbs up/down feedback UI enabled on all agent messages via feedbackLoop. User reactions and optional free-text comments captured to DynamoDB (staffinity-agent-feedback table) with 2-year TTL. Positive/negative/unknown classification. Original message content preserved via Redis cache. Data queryable for quality review and model improvement.

MG | MANAGE

Prioritizing and Addressing AI Risk

Responding to AI risks and improving AI systems over time. MANAGE activities address identified risks through documented processes, ongoing monitoring, and continuous improvement cycles.

Staffinity Controls Mapping

MG-1.1	Incident Response Plan documented (IRP v1.0, May 2026). Covers AI-specific failures including hallucination-caused harm, data leakage events, tool execution failures, and third-party model outages.
MG-2.2	Circuit breaker implemented in production: 3 consecutive tool/API failures trigger automatic service escalation and human notification via email + internal alerting. Agent halts until issue is acknowledged.
MG-2.4	Model update process: new Anthropic model versions evaluated in staging environment before any production deployment. Clients notified of significant model changes with 30-day advance notice where feasible.
MG-3.1	Continuous improvement: monthly usage metrics reviewed by Brent Sudeck. Client feedback incorporated into system prompt updates, tool configuration changes, and knowledge base refresh cycles.
MG-3.2	Vendor risk management: Anthropic assessed quarterly against security posture and service reliability. Sub-processor list published in DPA (trust.staffinity.io/dpa.html); 30-day advance notice of sub-processor changes.
MG-4.1	Decommissioning: upon contract termination, all client data deleted within 30 days per DPA Section 10. Deletion certification provided to client. KMS CMK scheduled for deletion after retention period expires.

CONTROL SUMMARY TABLE

The following table summarizes all NIST AI RMF control mappings documented in this alignment assessment.

Control ID	Function	Description	Status
GV-1.1	GOVERN	AI Governance Policy documented	Implemented
GV-1.2	GOVERN	AI Security Officer designated	Implemented
GV-1.3	GOVERN	Roles & responsibilities documented	Implemented
GV-2.1	GOVERN	Risk tolerance defined	Implemented
GV-4.1	GOVERN	Risk mgmt in deployment process	Implemented
GV-5.1	GOVERN	Per-client deployment review + DPA	Implemented
GV-6.1	GOVERN	Annual policy review schedule	Implemented
GV-7.1	GOVERN	Third-party AI provider evaluation	Implemented
MP-1.1	MAP	Use case scoping per client	Implemented
MP-2.1	MAP	Stakeholder identification	Implemented
MP-2.3	MAP	AI context & autonomy documentation	Implemented
MP-3.1	MAP	Negative impact assessment	Implemented
MP-4.1	MAP	Benefits documentation per client	Implemented
MP-5.1	MAP	Autonomy level categorization	Implemented
MS-1.1	MEASURE	Evaluation metrics defined	Implemented
MS-2.1	MEASURE	Pre-deployment test suite (Yeager)	Implemented
MS-2.5	MEASURE	Bias and fairness review	Implemented
MS-2.6	MEASURE	Explainability / tool transparency	Implemented
MS-3.1	MEASURE	Continuous performance monitoring	Implemented
MS-4.1	MEASURE	Incident audit trail (WORM)	Implemented
MS-4.2	MEASURE	User feedback loop	Implemented
MG-1.1	MANAGE	Incident Response Plan (IRP v1.0)	Implemented
MG-2.2	MANAGE	Circuit breaker + auto-escalation	Implemented
MG-2.4	MANAGE	Model update evaluation process	Implemented
MG-3.1	MANAGE	Monthly metrics review + improvement	Implemented
MG-3.2	MANAGE	Vendor risk management (Anthropic)	Implemented
MG-4.1	MANAGE	Decommissioning / data deletion	Implemented

COMMITMENT & CONTACT

Staffinity's alignment with the NIST AI Risk Management Framework reflects our belief that trustworthy AI requires proactive governance, not reactive compliance. We view the AI RMF not as a checkbox exercise, but as a practical tool for building AI systems our clients can trust.

This alignment document is reviewed annually and updated as the NIST AI RMF evolves, including as NIST publishes AI RMF playbooks for specific sectors and use cases. Staffinity will also update this document when material changes occur in our platform, governance structure, or third-party provider relationships.

Document Information

Document Title	NIST AI RMF — Staffinity Alignment
Version	1.0
Prepared	May 2026
Owner	Brent Sudeck, CEO & AI Privacy and Security Officer
Review Cycle	Annual (next review: May 2027)
Classification	Staffinity Confidential — may be shared with clients and prospects under NDA

Contact

Security & Compliance questions: **security@staffinity.io**
Trust center and documentation: **trust.staffinity.io**

> *"Responsible AI is not a destination — it's a discipline. We govern our AI systems the same way we govern our security posture: with documented controls, clear ownership, and continuous improvement." — Brent Sudeck, CEO, Staffinity, Inc.*