AI Usage Guidelines Package

Northstar SaaS Co. | Prepared with GuardAxis

Northstar SaaS Co. receives a high draft risk posture in this Phase 2 demo because the intake indicates internal and external AI use with customer content, source code, trade secrets, credentials or secrets considerations.

This package is a draft and requires business, security, and legal review before use.

At a Glance

- Included documents: AI Acceptable Use Policy; AI Security and Governance Policy; Employee AI Usage Standard; Third-Party AI Review Checklist.
- Evidence basis: 6 confirmed business inputs and 4 supporting website evidence notes.
- Drafting basis tracked for 8 sections before export.
- Priority risks carried into drafting: 3.

Executive Summary

- Intake answers are treated as the primary source of truth.
- Website evidence is limited to bounded public page review within the requested domain family.
- Policies are presented as editable draft guidelines.

Assumptions and Missing Information

Assumptions

- The website scan in this phase is bounded to public pages within the requested domain family and may be incomplete.
- Framework mappings are provided as high-level drafting references only.

Missing Information

- Public pages do not confirm the exact guardrails used for customer-facing AI outputs.

Package Snapshot

This draft guidelines package is grounded in 6 business inputs, 4 website evidence notes, and 3 priority risks. Detailed grounding appears in the appendices so the policy language itself stays readable.

Priority Risks

- Customer-facing AI outputs need clear human review boundaries (Medium)

- Source code and proprietary information exposure risk is elevated (High)
- Third-party AI tools need consistent approval and vendor review (Medium)

AI Acceptable Use Policy

AI Acceptable Use Policy

Audience: All staff and contractors | Status: Draft

Draft guidelines only. Review by business, security, and legal stakeholders is required before adoption.

Purpose and scope

Northstar SaaS Co. may use approved AI tools to support legitimate business work. These draft guidelines apply to employees and contractors and should be reviewed by business, security, and legal stakeholders before adoption.

Restricted and prohibited uses

AI may support approved work, but it must not be used for prohibited activities such as Do not submit credentials, secrets, or production source code into unapproved AI tools.; Do not use AI to make autonomous customer commitments or support decisions.; Do...

AI Security and Governance Policy

AI Security & Governance Policy

Audience: Leadership, security, IT, and governance owners | Status: Draft

Draft guidelines only. Review by business, security, and legal stakeholders is required before adoption.

Roles and responsibilities

Business owners, security reviewers, and legal stakeholders should share responsibility for approving AI use cases, reviewing exceptions, and confirming that policy language matches actual operations.

Logging and monitoring

AI-assisted workflows should follow the organization's logging expectations, with additional attention on customer-facing outputs and other high-impact uses. Incident review and exception tracking should be documented in a lightweight, repeatable way.

Employee AI Usage Standard

Employee AI Usage Standard

Audience: Employees and managers | Status: Draft

Draft guidelines only. Review by business, security, and legal stakeholders is required before adoption.

Daily use rules

Employees may use approved AI tools for Internal drafting and summarization, Engineering assistance and code review support, Support ticket triage and response drafting, Product and operations analysis. Customer-visible or externally shared outputs should s...

Review and verification

AI output must be checked for accuracy, completeness, and business appropriateness before it is relied on. When facts are uncertain or the impact is high, employees should escalate rather than assume the output is correct.

Third-Party AI Review Checklist

Third-Party AI Review Checklist

Audience: Security, procurement, and business owners | Status: Draft

Draft guidelines only. Review by business, security, and legal stakeholders is required before adoption.

Vendor basics

Record the tool name, business owner, intended use case, and whether the tool may affect customers directly. If the use case changes materially, repeat the review.

Data exposure and retention

Review what information the vendor receives, stores, or uses for model improvement. Pay special attention to customer content, source code, trade secrets, credentials or secrets, and confirm whether those categories are allowed at all.

Appendix A | Business Inputs

- Organization: Northstar SaaS Co.
- Primary website: https://software-saas.guardaxis.io
- AI usage mode: internal and external
- Customer-facing AI: In scope
- Sensitive data types: customer content, source code, trade secrets, credentials or secrets
- Current AI tools: ChatGPT, GitHub Copilot, AI features inside productivity tools

Appendix B | Observed Website Evidence

- software-saas.guardaxis.io: Northstar SaaS Co.

  Observed public fact | Source: software-saas.guardaxis.io

- software-saas.guardaxis.io/operations: Product teams use AI for internal drafting, engineering assistance, support ticket triage, and customer-facing workflow assistance with human review.

  Observed public fact | Source: software-saas.guardaxis.io/operations

- software-saas.guardaxis.io/trust: The sample trust page discusses vendor review, security review expectations, access control, and separation between customer data and internal productivity workflows.

  Observed public fact | Source: software-saas.guardaxis.io/trust

- software-saas.guardaxis.io/privacy: The sample privacy notes describe customer content, account data, support data, and restrictions around credentials, secrets, and production source code.

  Observed public fact | Source: software-saas.guardaxis.io/privacy

Appendix C | Priority Risks

- Customer-facing AI outputs need clear human review boundaries (Medium)
- Source code and proprietary information exposure risk is elevated (High)
- Third-party AI tools need consistent approval and vendor review (Medium)

Appendix D | Drafting Basis by Section

AI Acceptable Use Policy

Purpose and scope

- Business input: AI usage mode
- Business input: Customer type
- Drafting plan: Purpose and scope

AI Acceptable Use Policy

Restricted and prohibited uses

- Risk driver: Customer-facing AI outputs need clear human review boundaries
- Risk driver: Source code and proprietary information exposure risk is elevated
- Risk driver: Third-party AI tools need consistent approval and vendor review
- Drafting plan: Restricted and prohibited uses

AI Security and Governance Policy

Roles and responsibilities

- Business input: New tool approval
- Business input: Vendor review requirement
- Drafting plan: Roles and responsibilities

AI Security and Governance Policy

Logging and monitoring

- Business input: Logging requirement
- Risk driver: Customer-facing AI outputs need clear human review boundaries
- Drafting plan: Logging and monitoring

Employee AI Usage Standard

Daily use rules

- Business input: Internal AI use cases
- Drafting plan: Daily use rules

Employee AI Usage Standard

Review and verification

- Business input: Human review requirement
- Drafting plan: Review and verification

Third-Party AI Review Checklist

Vendor basics

- Business input: Current AI tools
- Drafting plan: Vendor basics

Third-Party AI Review Checklist

Data exposure and retention

- Business input: Sensitive data types
- Drafting plan: Data exposure and retention

Appendix E | Control Mapping

Framework mappings support review and drafting only and should not be treated as proof of compliance.

Review Disclaimer

Draft guidelines only. This material does not establish compliance, certification, or legal sufficiency and must be reviewed by business, security, and legal stakeholders.