

Governing AI-Generated Content at Scale

Paige Stevens · September 2026

AI writes content. Governance makes it trustworthy.

I'm part of the team defining a governance plan to keep AI-generated customer content accurate, consistent and compliant across a healthcare SaaS portfolio.

MY ROLE

Senior Technical Writer & Content Strategist. Core team member helping define a program-level AI content governance plan, being presented to leadership for implementation in 2027.

COMPANY

Enterprise healthcare SaaS provider (details changed to protect confidential information)

SCOPE

Release notes, help topics, onboarding guides and in-app assistance across dozens of products

PARTNERS

Technical writers, product managers, engineering, and AI and leadership stakeholders

SKILLS SHOWN

AI output evaluation

Prompt standards

Terminology systems

Content quality measurement

Content strategy

CONTEXT

A successful experiment raised a harder question

Our customers are hospitals and health systems. Their staff rely on our documentation to support more than 5 million healthcare workers, so incorrect instructions can cost them real time and trust.

In 2025 I helped launch AI tools that draft release notes and summarize help content. They worked: repetitive documentation time dropped and first-click success rose. That work earned a Gen AI Innovation Challenge Award.

If AI could draft release notes, help topics, onboarding guides and in-app help for every product team, who would make sure all of it was right?

THE PROBLEM

Small issues could quietly become systemic

The risk wasn't one bad help topic. It was thousands of AI-generated pieces across dozens of products.

RISK	WHAT IT LOOKS LIKE
Hallucination	AI invents a setting, workflow or limitation that doesn't exist.
Terminology drift	One product says "Facility," another "Location," another "Site" for the same thing.
Voice drift	Each team writes its own prompts, so every product sounds different.
Generic output	Content is technically correct but doesn't tell users what to do next.

No one owned these risks. Leadership was focused on how much AI could produce, not on how to keep it trustworthy.

APPROACH

A definition first, then two layers of controls

We started with a definition leadership could repeat:

AI governance for customer documentation is the set of standards, review processes, ownership models and quality controls that keep AI-generated content accurate, consistent, compliant, user-focused and aligned with the customer experience.

Working definition for the leadership presentation

Then we organized the controls into two layers. The first stops problems before content is drafted. The second assumes some will get through and catches them before customers do.

Layer 1

Prevent

Stop problems before content is drafted

- Approved sources only: no source, no content
- Citations back to source requirements
- Enterprise glossary and controlled vocabulary
- Centralized prompt library and templates
- Shared system instructions



Layer 2

Detect

Find what slips through before customers do

- QA review: style guide, voice, terminology
- Reading level, UI formatting, accessibility
- A risk score for every generated topic
- Scored on accuracy, compliance, consistency, completeness and customer focus

The plan uses two layers to catch AI content problems before customers do.

Three decisions shape the plan

No source, no content

AI may only draw from approved systems: Azure DevOps work items, requirements and approved help topics. This is the single strongest defense against hallucination.

Standardize the inputs, not just the outputs

A shared prompt library and templates prevent voice drift at the source, instead of fixing it topic by topic.

Measure, don't just review

A writer can't read every topic at scale. Scoring every topic lets people focus on the riskiest content first.

The plan shifts the technical content role from drafting every word to governing content standards, prompt standards, information architecture, quality measurement and advocating for our users.

THE WORK

Four pieces of the plan

01 A prompt standard

Every shared prompt follows the same rules, so any team gets consistent, source-grounded output.

ROLE	You write customer-facing release notes for a healthcare SaaS product.
SOURCES	Use ONLY the attached work items and approved help topics. If a detail is not in the sources, do not include it. Flag it as [NEEDS SOURCE] .
TERMINOLOGY	Use terms from the approved glossary. Never substitute synonyms.
FORMAT	One-line summary of the user benefit, then what changed, then any action the user needs to take.
VOICE	Plain language, casual but still professional. Active voice, second person. No marketing claims.

02 A terminology entry

One governed term, applied everywhere. This is where AI can reduce risk: a controlled vocabulary lets automated checks catch drift that humans miss.

PREFERRED TERM	DON'T USE	DEFINITION	NOTES
Facility	Location, Site	A physical place where an organization delivers care	Use "Facility" in UI labels, help and release notes. "Site" is reserved for the web address.

03 A quality rubric

Under the plan, every AI-generated topic is scored before it's published.

RISK AREA	QUESTION THE REVIEWER ASKS
Accuracy	Is every statement tied to source content?
Compliance	Could any statement conflict with healthcare rules?
Consistency	Does the terminology match the approved glossary?
Completeness	Are prerequisites and next steps explained?
Customer focus	Does it tell users why this matters and what to do?

04 Before and after

Illustrative example with a fictional feature.

RAW AI DRAFT

We're ~~excited~~ to announce ~~powerful~~ new enhancements to ~~Location~~ management! Admins can now ~~seamlessly~~ bulk-import ~~sites~~, and the new auto sync feature ~~keeps everything up to date in real time~~.

AFTER GOVERNANCE REVIEW

Add multiple facilities at once. Admins can now import up to 500 **facilities** from a CSV file instead of adding them one at a time. To start, go to Settings > Facilities > Import.

What changed and why

Removed "auto-sync"

It wasn't in the source work item. No source, no content.

"Location" and "sites" became "facilities"

Glossary term.

Cut "excited," "powerful" and "seamlessly"

Marketing words that don't help the user act.

Led with the benefit and added the path

Users now know what changed and what to do next.

OUTCOME

Where it stands

The plan is being presented to leadership now, with implementation planned for 2027. It gives leadership a shared way to talk about AI content risk and names the content team as its owner. It builds on AI documentation tools I helped launch in 2025:

30%

Less time on repetitive documentation with AI-generated release notes

40%

Improvement in first-click success from AI-powered content summarization

2025

Gen AI Innovation Challenge Award for the underlying AI work

WHAT I LEARNED

AI makes words cheap, which makes judgment more valuable. The most useful thing a writer can do at scale isn't write faster. It's to decide what "good" means, write it down, and build the system that holds every piece of content to it.

Paige Stevens · Senior Technical Writer & Content Strategist · Tucson, AZ