

Keep Content on Track With AI Guardrails for Content Standards

Create Better Content, Faster — Without the Risk



If you're like many Acrolinx customers and work in an enterprise, you probably have a large brand style guide. It's usually hundreds of pages long, designed to keep writers aligned with your company's style, voice, tone, and terminology. With enterprises now adding Large Language Models (LLMs) into their content supply chain, a style guide needs to not only inform writers, but also machines.

While we still write content for people, we also want to make sure our models can process it easily. With the addition of generative AI, there's a need for new and different ways to improve the process to ensure you're not compromising quality. In the era of AI, you need enterprise-wide content guardrails that protect your business — and your customer experience.

Ask yourself:

1. Do we have the technology and processes in place to ensure only high quality content is used for grounding our model?
2. Do we have the ability to control AI generated content before it is presented to a writer to make sure it meets corporate standards and uses compliant terminology?
3. Are we ready to quality control the large influx of content that will be generated before it's published without missing deadlines?

Did you answer yes to any of the questions above? Then you need AI content guardrails to maintain those standards.

How enterprises use AI content guardrails in their content supply chain

Content governance takes the key elements of your content strategy — your goals, priorities, and guidelines — and turns them into actionable guidance for writers and metrics that allow you to optimize your strategy. **Guardrails are one part of governance in action.** They're the content quality assurance steps in your content supply chain that might include:

→ **Quality check before fine-tuning:** Checking the current quality and attributes of the content you're using to train your LLM with before using it as data for fine-tuning.

→ **Quality check in workflows:** Integrating content quality checks into the human-powered creation process. Such as **automated quality gates** at different stages in the content supply chain, e.g. when content is published from a CCMS, to block non-compliant content from being published.

→ **Quality check on generated completions:** Checking that the content generated by LLMs meets your standards and expectations, before it gets to human readers. How well does it align with your **style guide** and where does it deviate? Measuring content quality at this stage is important to give feedback to your LLM about where it's generating inaccuracies.

→ **Quality check of published content:** **Auditing and analyzing** the quality of content that's already published, to ensure it stays up to date and aligned with your standards.

The ROI of having guardrails in place

Your writers are the experts in their field, and Acrolinx is the expert in your company's style guide. Acrolinx strikes the perfect balance by using AI to enhance creativity and prevent writer's block, while still ensuring that content aligns with your standards.

Our Average Customer Impact

+6M

Style guide violations identified per customer



+81%

Non-spelling and grammar



Identifying issues in compliance, inclusivity, clarity, terminology and more.

“These violations represent true business risk.”

With Acrolinx, you produce great and compliant content efficiently — saving you time and money.

Why our customer love Acrolinx for content governance:

1

Acrolinx's digital style guides help customers to tailor content governance according to their specific content types and writing standards.

2

Acrolinx integrates into your generative AI workflows to check, score, and improve content as it's generated.

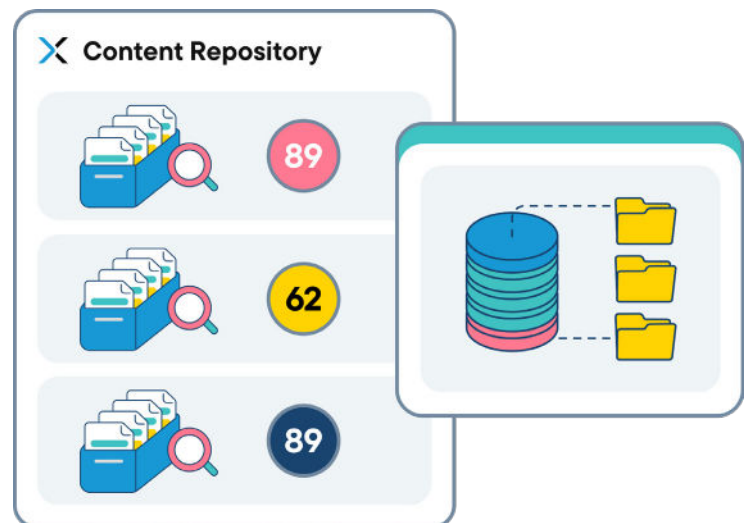
3

Acrolinx checks for clarity, tone, and inclusive language. It also helps customers scale content workflows and leverage generative AI without compromising on quality.

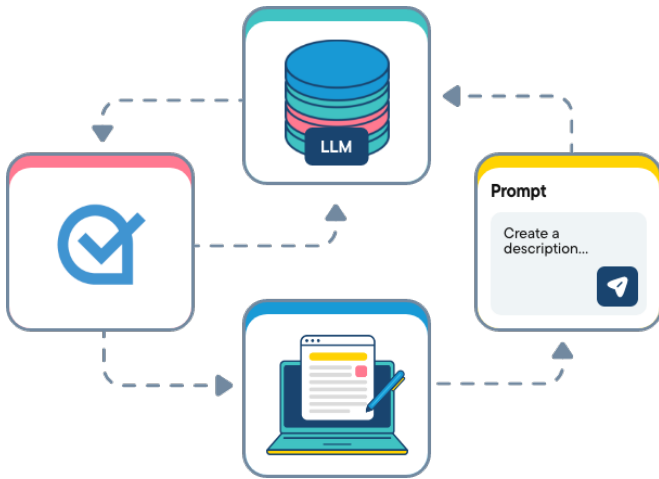
Governance and guardrails use cases

Fine-tune your LLM with quality assured content

Your LLM can only produce quality output when fine-tuned with quality source content. So how do you know which of your content assets will lead to clear, consistent, and on-brand AI-generated text? Acrolinx can check batches of your existing content and help you make business decisions about what to use to fine-tune your LLMs. This drastically improves your model's output.



Check your LLM output for brand alignment



How do you validate your LLM-generated content is on-brand? Or prevent “drift” in the model — where the output becomes less accurate over time? Start by checking the quality of the content your LLM generates — before your writers even see it. Acrolinx integrates into your generative AI workflows to check, score, and improve content as it’s generated. It makes sure content follows enterprise style guides and writing standards, by checking and scoring how well it matches your expectations.

Deploy enterprise-ready generative AI fast

Expedite content creation that meets your standards with the Acrolinx AI Assistant and Get Suggestions. These new generative AI capabilities deliver instant value to enterprise writers. Drastically reduce content errors and editing cycles with easy “click-to-fix” content improvements that align content to your enterprise writing standards. The AI Assistant generates new content that aligns to those standards, while the Sidebar provides click-to-replace suggestions for 87% of writing issues. Working with the Microsoft Open AI Service, Acrolinx blocks the generation of offensive, abusive, hateful, and explicit content.

The image displays a document snippet on the left and the Acrolinx web interface on the right. The document snippet includes a red Acrolinx logo and a section titled 'Tone of Voice' with text: 'Content with the tone of voice that makes the audience go “Oh yeah, that sounds like something I’d like to have.” It balances a lively tone of voice with essential information your audience needs'. The web interface shows a 'Check' button, a status '334 words and 4 issues', and a list of suggestions: 'Here is Here's' (with a note 'Could you make it sound less formal?'), 'Acrolinx Acrolinx' (with a note 'Spelling'), 'actual fact fact' (with a note 'Could you make it simpler?'), and 'insanely' (with a note 'Is your language respectful to other...'). A 'Time Range' filter is set to '02/03/2023 - 02/03/2024', and there are dropdowns for 'Filter by Language' and 'Filter by Writing Guide'. A '3652 checked' indicator is visible at the bottom right of the interface.

Ready to safeguard your content standards with Acrolinx?



Better Content. Faster.

Acrolinx is your enterprise content insurance policy. Our AI-powered content governance software captures and digitizes your style guide to make your writing standards, standard.

Acrolinx governs new and existing content written by people and generative AI. Whether your company has written 100,000 words or billions (like our customer Microsoft), Acrolinx makes sure each one reflects your style guide. Customers enjoy massive efficiency gains without sacrificing standards through AI-powered live writing assistance, automated reviews and quality gates, and analytics comparing content quality with performance.

With an LLM infrastructure anchored in Azure AI, Acrolinx guarantees scale, future-readiness, and uncompromising safety and security. Born out of the German Research Center for Artificial Intelligence (DFKI), AI runs deep in Acrolinx's DNA.

www.acrolinx.com

Acrolinx® is a trademark of Acrolinx GmbH.

Copyright © 2024, Acrolinx GmbH, all rights reserved.