


Building the Future

Real-world insights into generative AI
adoption and risks in the enterprise

An Acrolinx report



The dawn of generative AI

Everyone's talking about generative AI, but how many enterprises are actually using it? And if so, what types of content do they plan on using generative AI for?


As generative AI technology disrupts the content creation space, enterprises are considering what role it should play in their content strategy. But with public opinions and approaches fluctuating daily, it's time to evaluate what enterprises are doing on the ground.

In this report, we set out to understand if enterprises are using generative AI and why, what concerns them the most about it, and how they think generative AI models should be trained. Our survey was taken by a mix of Acrolinx Fortune 500 customers (including Amazon, HPE, Humana, Fiserv, and Intralox) and other enterprises and received 86 responses.

Let's explore the findings.



Are enterprises using generative AI to create content?



Overwhelmingly, **95.3 percent of respondents are either already using generative AI to create content or planning to do so.** Compared with the 4.7 percent of respondents who aren't using it, it's telling how many enterprises want to embrace generative AI.

Suffice to say, enterprises are taking steps to substantiate all the chatter around generative AI. The vast majority of them either already have or are planning to integrate generative AI into their content creation workflows.

If not, why not?

Before we move on, let's dive into the 4.7 percent of respondents who aren't using and don't plan on using generative AI.

Out of 86 respondents, only four claimed they weren't using generative AI. But why?


- **“Not sure how to implement [it] or how it would work.”**
- **“Our content is too varied and requires thorough fact-checking. Generative AI has a well-earned reputation for being confidently wrong.”**
- **“Not [planning to] yet.”**
- **“We're considering AI to help check quality, but not author.”**

The main themes emerging from the survey center around uncertainty for how best to use generative AI and concerns over the accuracy of AI generated content. These themes align with the public skepticism about generative AI.

As the generative AI field evolves, best practices and standards will develop and help alleviate hesitancy. However, the generative AI model that enterprises choose to use has implications for accuracy — but more on that later.

Would you use it if your worries were addressed?

The four respondents that said they aren't using generative AI to create content were also asked if their issues with using it were addressed, would they use it.



With a 50/50 response, it's clear that those who aren't planning to use generative are divided on how open they are to using it in the future, so long as their concerns are alleviated.

With so few respondents saying they don't use generative AI, it leaves only a small sample to answer this question. But interestingly, even those 50 percent currently opposed to using generative AI would change their stance if their issues were addressed. **This confidently demonstrates that generative AI is taking up an ever-increasing foothold in the content landscape.**

Within the responses, several themes emerged in terms of the goals enterprises hope to achieve with generative AI.

- **Efficiency.** Lots of answers centered around driving more efficiency, creating more content in less time, and reducing costs. Some answers emphasized achieving this through generative AI creating first drafts and using more automation.
- **Improve quality.** Enterprises want to use generative AI to create better content, stimulate creativity, create more summarized content, and make sure all content is more consistent.
- **Technical and support content.** Some answers highlighted how generative AI could create easier-to-read documentation, generate DITA content, develop “emotionally intelligent” chat responses, and generally help deflect support requests.
- **Process improvement.** A few answers focused on how generative AI could improve content processes. For example, making sure language translation is consistent and professional and improving content management and administration.



The majority of responses were five words or less. We've included some examples from longer-form answers:


- “Improve [the] quality and comprehensiveness of what we write. Use it for editing tasks. See what others have said on topics of interest.”
- “My goal is to assist and provide helpful responses to users in various tasks, such as answering questions, providing recommendations, and completing tasks. I strive to continuously improve my language processing abilities and provide accurate and relevant information to users. Ultimately, my objective is to enhance the user’s experience and make their interactions with me seamless and efficient.”
- “Use it as a quality check, but not a content author.”
- “I would like to make it possible to have templates for many types of content, in order to help the source content teams produce their use cases. On top of that, it would be great to have AI consider what is available in certain target markets, and leverage that knowledge for localization.”
- “Generative AI is good at creating “boilerplate” documents as a starting point for a human. In our particular organization, I don’t think that creating customer help pages would benefit from this process any more than starting from existing content.”

What’s evident in the goal setting of enterprises is an awareness of generative AI’s limitations. Many responses highlight how AI generated content can be used as a template for humans to build upon or alternatively as a quality and consistency checker. Likewise, generative AI is seen as a starting source of inspiration for content creators.

Needless to say, **enterprise goals for generative AI content creation align to previous content creation goals: greater productivity, stronger brand alignment, streamlined processes, and general compliance.** As such, generative AI is perceived to help enterprises meet their current goals and is seen as part of the content workflow.

What types of content would you use generative AI to create?

We asked our respondents to identify the types of content they'd use generative AI to help create.



Marketing content is narrowly the most popular at 67.4 percent, with product content 1.1 percent behind at 66.3 percent. 60.5 percent of respondents would use generative AI for technical documentation content and only 53.5 percent for support articles.


With just over 12 percent variance from the least to most popular content type and all answers being higher than 53 percent, it's fair to say that those that want to use generative AI are willing to use it across lots of content types. Although at almost 21 percent, technical documentation was cited as the content type respondents would be least likely to use generative AI to create.

Interestingly, although several respondents highlighted using generative AI for support content, it was the least popular content type. But considering that enterprises want to use generative AI as a creativity stimulus, it's not surprising that marketing content was the most popular choice.



Does your organization currently restrict the use of generative AI?

We're curious about how many enterprises restrict the use of generative AI by their employees.





The results show that an equal percentage of respondents said their enterprise restricts versus doesn't restrict generative AI — 41.9 percent. It's surprising that **a staggering 41.9 percent of enterprises have no restriction on the use of generative AI.**

Notably, 16.3 percent are unsure if their enterprise restricts generative AI use. It's easy to observe that without clearly communicated use guidelines, those respondents who are "unsure" could be using generative AI, but their organization might not sanction it.

Do security concerns influence your decision to use generative AI?

Much of the discussion around generative AI has centered around security. We felt it's important to ask whether security concerns influence enterprise decisions when it comes to generative AI.



The majority of respondents (68.6 percent) said that security concerns have influenced their decision. 9.3 percent were unsure if it had influenced their enterprise's decision, while 22.1 percent claimed it didn't impact their decision to use generative AI. This is in keeping with current discourse, as the majority cite it as a concern.

What's your biggest concern?

As a follow up to our security question, we wanted to dive deeper into the biggest concerns of enterprises. Of the 86 respondents, 25 selected intellectual property as their biggest concern. This was followed closely

#1	Intellectual property concerns	#2.27 average
#2	Customer security compliance concerns	#2.52 average
#3	Bias and inaccuracy	#3.28 average
#4	Privacy risk	#3.43 average
#5	Public data availability and quality	#3.5 average


by customer security compliance concerns, which received 23 votes. Comparatively, the other three answers (bias and inaccuracy, privacy risk, and public data availability and quality) received between nine to 16 votes each.

The concern that was ranked the lowest by most respondents, at 43 votes, is privacy risk. However, as more people identified it as the biggest concern over public data availability and quality, it has a higher average.

It's worth noting that there's only a 1.23 point difference between the highest and lowest result. Without a clear largest concern, **it's fair to say that all the listed concerns are part of the decision-making process for enterprises.**

How should your generative AI model be trained?

With a lot of discussions around how tools like ChatGPT get their information, we wanted to find out how enterprises want to train their generative AI model.



Surprisingly, **the majority of respondents, 41.9 percent, would rather use an AI model that trains on publicly available content.** Just over a quarter of respondents would like to train their model on just their own content. And likewise just over a quarter of respondents would like a mix of both their own content and publicly available content, while only a small percentage (3.5) were unsure.

The preference for public content speaks to the vast availability of digital content that generative AI tools can use to train their models. Though it remains surprising that enterprises aren't more focused on using their own content to build models.

So where do we stand with generative AI?

With the vast majority of respondents either using or planning to use generative AI, it's clear that generative AI is already an established tool that enterprises plan to use.

When it comes to why they want to use it, enterprise goals are based around efficiency, improving content quality, generating technical and support content, and improving processes. There's a slight preference for developing marketing content with generative AI, but enterprises are open to using it to generate all types of business content.

In terms of concerns, almost 69 percent of respondents see security as an issue, despite the fact that only 41.9 percent of enterprises currently restrict generative AI use. Intellectual property came out on top for what concerns enterprises the most, which juxtaposes with the enterprise preference to train AI models on publicly available content.



As much as enterprises are turning to generative AI as a tool for content creation efficiency, it remains an emerging technology that enterprises are grappling to understand.

That's why Acrolinx created [AI Enrich](#) — so enterprises can enjoy the benefits of generative AI without the associated risk.

A safeguard against generative AI risk is content governance. And we just happen to be content governance experts. We're combining content governance with the power of generative AI. Now, with our first-of-its-kind enterprise-grade solution, enterprises can confidently and safely use generative AI as part of their larger content supply chain.

With Acrolinx, any content, whether generated by humans or through generative AI, is aligned with enterprise strategy, analyzed for performance by audience and objective, and continuously improved for the highest enterprise impact.



Some examples of answers include:

- “We currently use Acrolinx during content creation, to check and improve quality.”
- “Acrolinx helps me improve the quality of my documentation, suggest changes, and improve the accuracy of my language.”
- “Improve content quality and reduce costs.”
- “Sidebar integrations during [the] authoring of new content.”
- “We use it increasingly for consistency and quality in Technical Publications.”
- “Used for automated content checking.”
- “We use Acrolinx for all our content written in English. We now want to take it to the next level, and leverage the new AI feature.”
- “Ensuring associates use consistent terminology and follow authoring standards.”
- “Creating AI that works with, not instead of, people.”
- “Acrolinx already does a great job of being a highly customizable style and grammar tool.”
- “Consistency, clarity, and overall high quality content.”
- “Inclusivity.”

But now with AI Enrich we can go one step further to help enterprises. AI Enrich for Acrolinx allows your business to tune AI on your content — allowing employees to create new copy, generate texts, or improve its performance.

With an enterprise-grade solution, your organization can encourage and accelerate experimentation with generative AI, by governing which content collections to use to continually tune your AI model.

Better yet, Acrolinx can identify your best performing content, determine its quality standards, and use that content as tuning data for your AI model. Meaning you create more content that converts, informs, and engages your audience, based on your best content characteristics. Now you can sleep easy knowing that your employees are using generative AI to create content.

[Learn more](#)



About Acrolinx

Acrolinx is an AI-powered software that improves the quality and effectiveness of enterprise content. Acrolinx customers increase their content's value by streamlining its creation, governing it against established writing guidelines, and improving its impact over time. Acrolinx helps Fortune 2000 companies, like IBM, Google, and Microsoft, eliminate editorial bottlenecks, quality issues, budget overruns, and compliance risks from their content supply chain.

Acrolinx is a trademark of Acrolinx GmbH. © 2023 Acrolinx GmbH.