

Report: 2024 Generation Artificial Intelligence (GenAI) Benchmark Study

By the Center for Information-Development
Management (CIDM)

Table of Contents

- Summary.....2**
- Intro..... 3**
- Methodology..... 3**
 - Survey Design3**
 - Demographic Questions3
 - Foundational Questions.....3
 - Adopted Questions3
 - Investigating GenAI Adoption Questions.....3
 - Not Currently Investigating Questions3
- Demographics 3**
 - Represented Industries.....3**
 - Implementation status4**
 - Leading efforts.....4**
 - Sources of information8**
 - Information on GenAI in General8
 - Information on evaluating GenAI9
- Why GenAI..... 11**
 - Benefits..... 11**
 - Promising Applications for GenAI 13**
 - Content types that will benefit the most 16**
 - Challenges to implementing AI..... 18**
- Why Not GenAI..... 21**
 - Reasons for not implementing 21**
 - What would change the decision 22**
- GenAI and Content Creation 24**
 - How has it changed the content creation process..... 24**
 - Factors for why it hasn't changed the process 27**
- Evaluating GenAI 28**
 - As a tool..... 28**
 - Ethical Considerations..... 30**
- Trends 31**
- Conclusions 32**

Executive Summary

GenAI is not just a new tool—it’s a shift in how content is created and delivered. Getting the most out of this technology requires investment in skills, systems, and governance. This will decrease your documentation time to market, to protect your company’s data, and ensure quality. Supporting technical communication teams in this transition is key to staying competitive and meeting evolving customer expectations. This survey explores how technical communicators research, implement, test, and use GenAI in practice to increase capacity, improve quality, and limit rote tasks to focus on deep knowledge work such as critical thinking, content strategy, and documentation development. By leading and consulting on GenAI implementations, technical communicators are becoming strategic leaders in AI integration, moving beyond “just writing” to managing AI-driven processes and delivering AI-ready content.

Important Insights

- **Widespread Impact:** Nearly 80% of professionals expect GenAI to significantly transform how technical communication is done in the next 5–10 years.
- **Adoption Status:** 86% have already adopted or are exploring GenAI, with many technical communication teams leading these efforts within their organizations.
- **Humans remain in the loop:** GenAI hasn’t replaced existing content creation steps but helps teams work more efficiently and improve quality with human oversight remaining crucial to verify AI-generated content.
- **Leadership support:** Organizations with leadership support for GenAI adoption tend to implement more successfully.

Top Benefits of GenAI

- Faster content creation and improved quality
- Increased capacity, allowing teams to focus on higher-value tasks
- Cost avoidance by scaling output without adding staff
- Improved customer engagement and support through AI-powered chatbots

Key Applications Being Used

- Content analysis and optimization (e.g., summarizing, editing)
- Automated translation to reduce costs
- Chatbots that answer questions with greater flexibility than traditional scripted bots
- Predictive analytics to help plan content strategy

Challenges to Watch

- Integrating GenAI with existing systems can be complex
- Shortage of skilled personnel slows progress
- Data security and intellectual property concerns, especially with public AI tools
- High upfront costs and uncertainty about scaling effectively

Introduction

Generative artificial intelligence broke through to the masses in 2022 with the promise of revolutionizing many aspects of business, art, and literature. This survey, conducted by [The Center for Information Development Management \(CIDM\)](#), seeks to understand the impact of GenAI in the technical communication industry two years into this revolution.

Participants answered questions in four sections of content: About you, GenAI Fundamentals, and a section of questions based on the participants current use of GenAI (Adopted, investigating GenAI for adoption, and no current plans to adopt GenAI), and some closing demographic information. All participants were asked questions in the first two sections, and the third section was determined by the individual participants themselves.

The About You section comprised three questions about the participants' location, CIDM membership status, and their role in the GenAI implementation at their organization. The ten-question GenAI Fundamentals section sought to establish an understanding about the participants' base knowledge about GenAI. The last question of this section asked participants to self-identify their department's current GenAI status. This question was used to filter participants into questions specific to their implementation status.

Demographics

Represented Industries

The survey participants (107 in total) represent a wide cross-section of the technical communications industry covering software (47%), machinery (11%), consulting (9%), hardware (5%), as well as several other industries (see Fig. 1). Additionally, participants are selling products or providing services relatively evenly around the world (see Fig. 2).

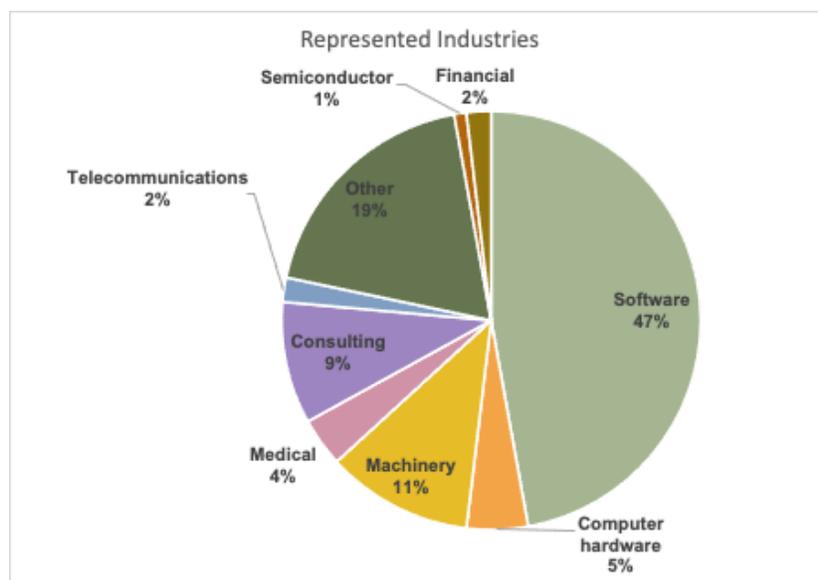


Figure 1: Represented industries

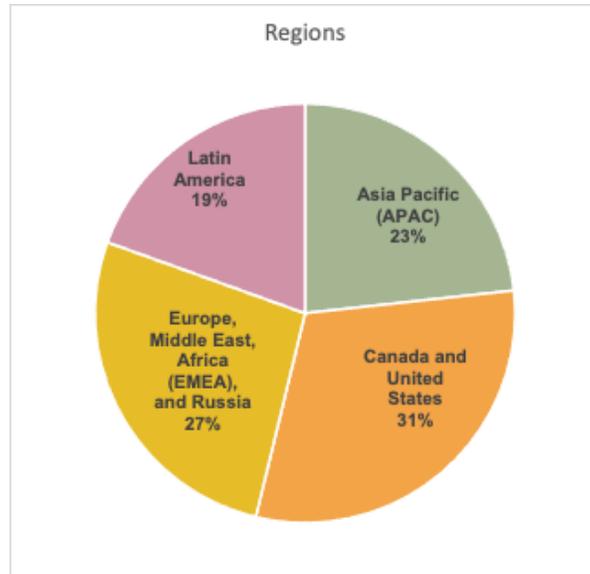


Figure 2: Sales regions

Implementation status

Our survey revealed that 86% of technical communication professionals have either adopted or are investigating GenAI for applications in the field. 29% of the participants are considered early adopters because they were able to adopt and implement GenAI within a year of it becoming publicly available.

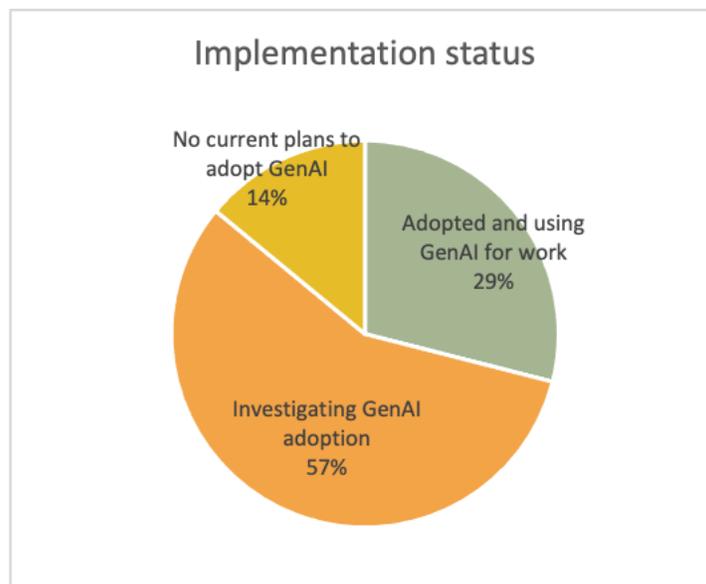


Figure 3: GenAI implementation status

Leading efforts

When we looked at the roles participants took in GenAI efforts, we saw that 69% of participants are involved in an influential role (see Fig. 4). Additionally, among the technical communication

professionals who have adopted GenAI, we saw that 87% are in influential roles with 35% championing GenAI efforts.

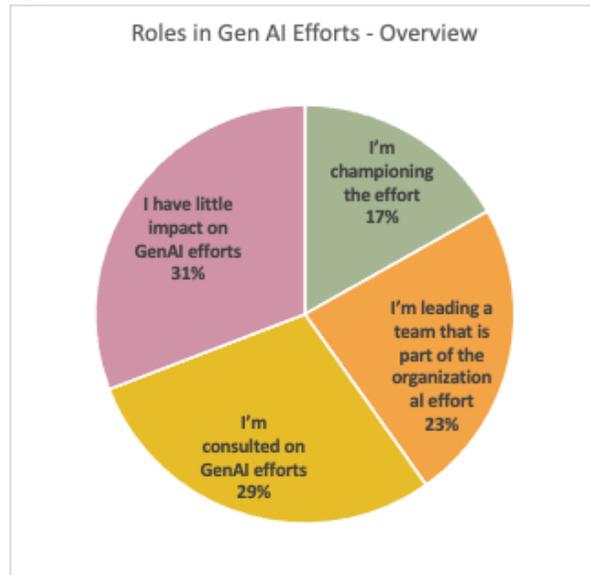


Figure 4: Roles in GenAI efforts

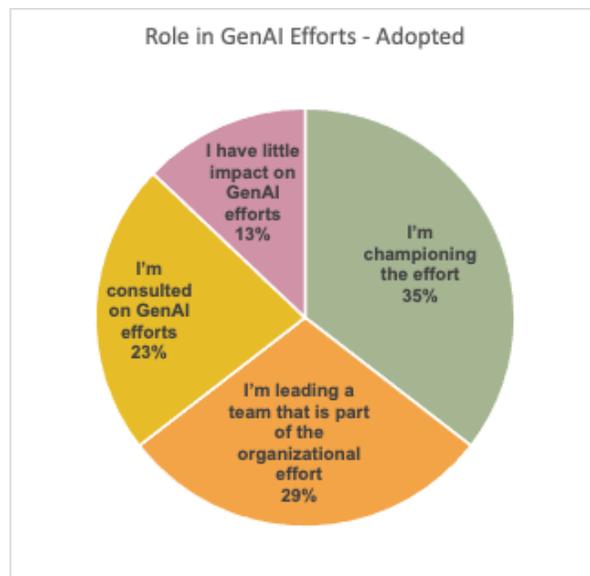


Figure 5: Roles in GenAI efforts - Adopted

29% of participants are leading teams that are participating in organizational efforts, while 23% are consulted. In sum, for organizations that adopted GenAI early, technical communications teams are heavily involved in 86% of the cases. Of that 86%, nearly half of the technical communication professionals (44%) are in highly influential roles in those implementations. (See Fig. 5). We're seeing that early GenAI efforts are being driven by technical communications teams and consultants. Speculatively, tech comm teams that are slow to adopt might be slowing adoption for their enterprises. However, we also note that many participants in the not adopting groups are in highly regulated fields and pending greater direction from various regulatory bodies. For example, medical, financial, and semiconductors have no champions among our participants. Telecommunications and consulting industries see the largest percentages of champions (50% and 60% respectively).

In the groups that have not yet or are not planning to implement we see slightly smaller percentages of champions (8% and 13% respectively) (See Fig. 6-7). Concurrently, we also see much greater percentages of technical communication professionals that have little impact on GenAI efforts, (38% for investigating and 40% for not interested in implementing). This change can be read as the organization is not interested or unable to implement GenAI and therefore technical communications is not needed or as technical communication is not the part of the organization where GenAI is going to be implemented. However, even in organizations that are investigating or are not interested in implementing, approximately 60% of our survey participants are at least minimally involved in GenAI projects.

This kind of heavy involvement in GenAI activities is one of the first indications that the field of technical communications is changing with GenAI. It highlights the fact that technical communication professionals will need to stay up-to-date on GenAI technology because they are likely to at least be consulted about GenAI and their organizations. As the technology and regulations evolve, more organizations may be able to implement GenAI, and it would be beneficial of technical communicators to be prepared for a change in direction.

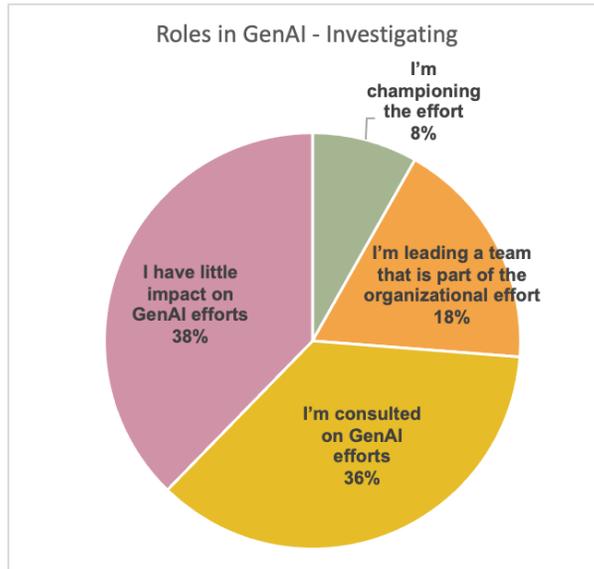


Figure 6: Roles GenAI efforts - investigating

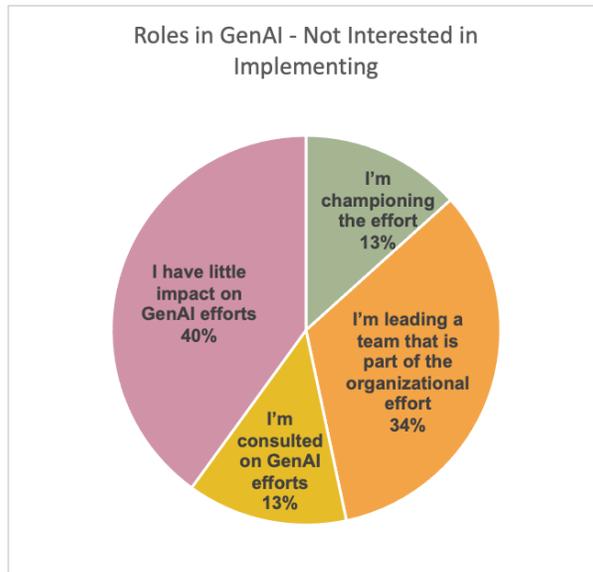


Figure 7: Roles GenAI efforts - not implementing

When the roles in GenAI efforts are examined by industry, we saw that the consultants, by far, have the largest share championing within their sector (see Fig. 8). Conversely, technical communication professionals working in the software industry have one of the smallest shares of championing efforts (see Fig. 8). However, the software industry had one of the highest percentages of team leaders participating in organizational GenAI efforts (see Fig. 8). Within the rest of the industries, technical communication professionals maintained large shares of influential roles for GenAI implementations, such as leading teams and championing efforts.

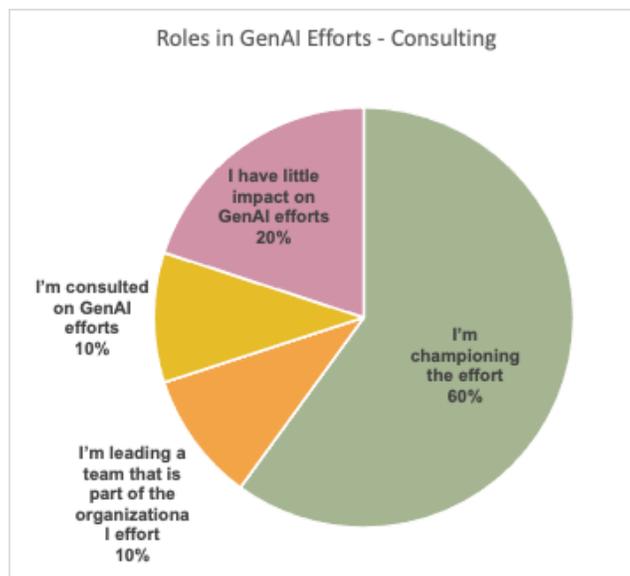


Figure 8: Roles in GenAI efforts - Consultants

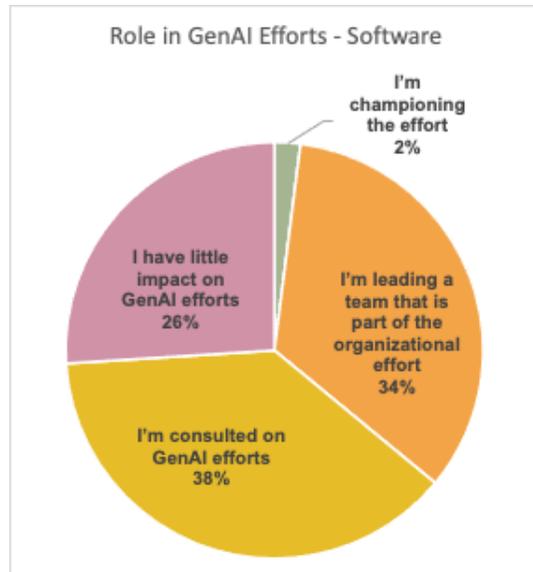


Figure 9: Role in GenAI efforts - Software

Sources of information

Information on GenAI in General

Given the high percentage of technical communicators championing, leading, and consulting on GenAI efforts, we looked closely at where they are getting their information. We saw that expert webinars and conference presentations topped the ranks as a very close one and two. Webinars and conferences leading the ranking indicate that technical communicators are eager for spoken and recorded information. (See fig. 10)

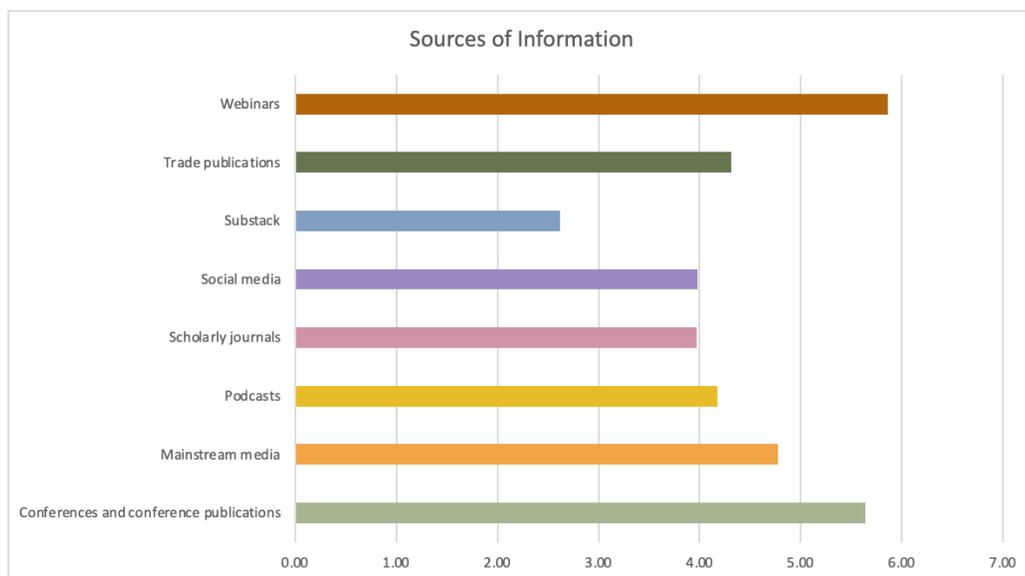


Figure 10: Sources of information

Webinars ranking at the top of the available options was unsurprising. Several key reasons for this are that they:

- provide nearly bite-size content in hour long blocks
- are readily available for minimal financial investment
- are usually asynchronous in that recordings can be provided for participants that cannot attend the live version and rewatched by folks who might have been distracted
- feature many different GenAI topics, implementation phases, products, viewpoints, and stories

One point of interest here, when looking for general information about GenAI, professionals are using mainstream media (ranked 3rd) before trade publications (ranked 4th) or scholarly articles (ranked 6th). We extrapolate that this was because they are following the broader implications, uses, and news (technology developments, laws and regulations, governance, etc.) about GenAI that comes from mainstream media because it is not limited to the technical communication industry; rather it likely includes the industry(s) they are writing for. However, they trust industry experts to tell them which GenAI tools, techniques, and approaches they should be using. They are looking for the technical communication industry information in spoken and discussion formats and they are not turning to the deeper written forms such as the trade publications and scholarly articles that would have the format and space to provide depth of information and research. (See Fig. 10)

Podcasts and trade publications are not statistically different to distinguish between fourth and fifth ranking (See Fig. 10). Podcasts and trade publications being closely ranked was likely because podcasts have evolved into a form of trade publication. There are several well-respected podcasts on technical communication, hosted by experts in the field. The content of these podcasts, while different in format, was comparable in substance to what these experts might be writing for the trade publications. Podcasts have the additional benefit of bringing guest speakers on which provides dialogue, multiple perspectives on a single topic, and a portable audio format.

Social media and scholarly journals are also not statistically distinguishable between sixth and seventh rankings (See Fig. 10). We infer that the level of content for both of them was disproportionate to what technical communication professionals are looking for. With social media the level of content and reliability was far too low to be considered reliable. At the other end of the spectrum are scholarly journal articles which tend to be far too dense and theoretical for industry professionals to be able to engage with the content on a regular basis. Additionally, social media, like mainstream media, was unlikely to be presenting GenAI information specifically for technical communications and scholarly articles are likely presenting information on all too specific scenarios. Both require high cognitive loads to apply any information from the source to an industry situation.

Unsurprisingly, Substack was consulted the least (See Fig. 10). Substack is a relatively new media outlet and not likely to be as widely adopted as several of the others. Additionally, Substack is a single application whereas some of the other choices of information sources are more broadly defined. Substack was included as its own line item in the list because it does offer access to many different authors, contexts, and information sources.

Information on evaluating GenAI

We asked participants who indicated they adopted GenAI where they went for information on evaluating GenAI. We were looking to understand where people who are working with GenAI are

getting their information for implementation. Among participants who adopted GenAI, expert webinars are the most popular sources of information with case studies being a close second. The third source for evaluating GenAI was conference presentations.

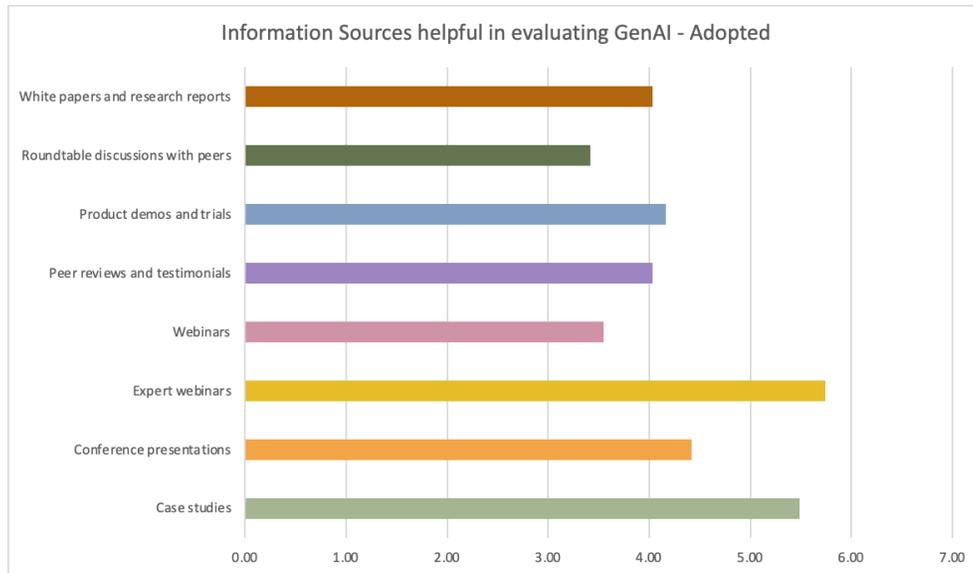


Figure 11: Sources of information to evaluate GenAI

Read together, we inferred that what early adopters find the most useful sources of information are case studies of implementations presented by experts in a spoken and demonstrative medium. There are many ways that GenAI can be used in technical communications to solve business problems, so people are looking to the broader industry to understand how these tools and technologies can be used, for what purposes, and how to implement them based on a specific use case or purpose.

It was important to note that the survey called out “case studies” and not “customer success stories”. While the two terms might be colloquially interchangeable, we presented the category as case studies to include unsuccessful, problematic, and/or incomplete implementations. Technical communication professionals will also benefit from understanding the challenges that others faced or are facing.

In interviews, the desire for case studies as a main source of information was confirmed. Interviewees also expressed a desire for a forum where they could actively participate in discussion with other professionals in similar situations. While they are content to view successful customer stories, they are asking for space to ask more questions and receive robust feedback. We theorize that while they are seeing customer success stories and tools implementations, there are missing links about how the entire system does or can look like in practice.

Even with all these information sources, many interviewees indicated that they were using a ‘trial and error’ approach. This appears to be an industry gap in education and information. In part, we suspect because the technology is so new, that there are very few customer success stories and few completed projects, meaning that the industry should be looking for some good customer success stories in the next year. Another potential reason was that creating a customer success story about a private large language model implementation without exposing proprietary and intellectual

property might be challenging given the amount of training on such content that was required to make the system work.

In contrast, peer reviews and testimonials ranked low on the list of information sources (6th of 8 sources). We understand this to mean that professionals are not turning to their peer networks for information but to trusted and verified sources to identify experts, such, a conference put on by a professional organization or a webinar hosted by a recognized industry expert.

In a similarly contrasting situation, interviewees (comprised on both early adopters and those looking to adopt GenAI in the next few years) indicated that they were interested in being put into conversation with other professionals working through GenAI implementation; however, among the early adopters, peer round tables were indicated as the least used resource for information. One likely explanation for this seemingly contradictory situation was that there were no or few roundtables available for the early adopters to participate in. A second explanation was that there are just so few early adopters (29% of our participants) that there are simply not enough “peers” to reach out to. Both explanations are understandable given how new the GenAI technology is in the technical communication industry; at least in terms of it being publicly and readily available.

Why GenAI

Benefits

GenAI has many potential and promising benefits. Whether looking at the aggregate data or at responses from the individual groups, the benefits of GenAI remain relatively constant. The greatest recognized benefit was the increase in efficiency when it comes to content creation (31% of all participants). The second recognized benefit was quality improvement (19% of all participants). Making up approximately 50% of the benefits, creating better content faster was the main benefit for implementing GenAI in technical communications.

Workload reduction (17% of all participants) was not surprising as the third benefit for two reasons. First, it is a more general category and has not had the same amount of media time and space as the first two benefits. Second, if quality and productivity increase, it should result in an increase in resources time to do other projects; effectively giving the appearance that workload goes down. However, in interviews, we saw that workload increased because additional projects or work could be taken on without adding head count.

Interviews emphasized that, while efficiency was gained, there was no reduction in the steps of the content creation process. The framework of getting information from subject matter experts (SME), drafting the content, reviewing with the SME, editing for grammar and style, and then publishing all remained intact. GenAI was used in a variety of these high-level steps, but a human remained in the loop at each stage to verify and amend the content that AI generated. Capacity and quality increased because humans were able to concentrate on tasks with higher knowledge-work and critical thinking. See GenAI and Content Creation for further discussion.

Cost savings was the fourth benefit (13% of participants). Initially, this was believed to be driven by things like decreases in time to create content. However, interviews with people who have already adopted GenAI indicated that cost avoidance, rather than cost savings, was occurring. “Cost savings” implies that there will be an immediate decrease in the cost of producing content.

This was simply not true because in many cases GenAI was added to the tool stack that teams are using and all the people required remained on the teams at full or nearly full capacity, i.e. there was no reduction in their job responsibilities. What this means was that GenAI and custom private LLMs create an additional cost for the enterprise. However, because the current staff can produce better content faster, there was cost avoidance of either having to add people in the future, or add other software (editing tools, for example) later. GenAI was helping teams scale the amount of content they can successfully produce without having to expand their resources again in the future; this was “cost avoidance”.

Enhanced customer engagement came in as the fifth recognized benefit (12% of all participants). Given that customer engagement is one of the foundational considerations of technical communication, it was surprising that this came in so low; especially when virtual chat assistants (VCAs) are one of the great new content delivery mechanisms. We suspect that seeing GenAI as a content delivery mechanism is a secondary understanding of the technology, particularly when people are thinking about the public LLMs, such as ChatGPT, and associated privacy and security risks. However, it is short sighted to not consider these public LLMs as forms of customer engagement; we are seeing them used on sites like Google and being used for research. Additionally, the use of private LLMs to power VCAs and chatbots is going to create another opportunity for targeted content delivery that engages customers in product or simply before they call support.

A common response to the open-ended Other category was that there was no benefit or that a benefit had yet to be realized (54% of Other). Many participants of these responses self-identified as being “somewhat familiar”, “somewhat unfamiliar, and “very unfamiliar” with GenAI. At least one of them also stated that they are not currently using GenAI. The one participant who self-identified as “very familiar” with GenAI stated “I perceive no benefits in development of original information”. This participant makes an important distinction between “development of original information” and all the other kinds of content that can be generated with AI. In fact, this was quite true and very nuanced because GenAI creates content by making word predictions based on statistical analysis of word use in language. Therefore if information was truly new, then GenAI’s statistical prediction could be uninformed and unhelpful.

Another common set of responses to the Other category were around consistency (15% of Other). These responses identified “consistency” and “uniform style and tone across departments”. Generative AI can be trained to a specific style guide, including tone and voice of writing. It was surprising that this response was not called out by more people, however, it can easily be considered part of the improving quality category. Many folks are seeing this kind of benefit with grammar tools that are not traditionally considered generative AI, but do offer suggestions for improving the style or grammar of the writing.

Several responses to the Other category indicate that GenAI could be used to self-edit particularly for non-native English speakers and writers. (15% of Other) This was an important highlight because it potentially fits into both main benefits of increasing efficiency and improving quality. It also points to translation as a benefit that was not specifically listed in our survey. The efficacy of GenAI for translation is still being debated and faces challenges in reviewing for technical communication teams that do not have an in-house translator or are otherwise not translation

specialists. However, the open responses indicate that GenAI is being used as an assistant to provide drafts for content creators not working in their primary language; the same as it was for content creators who are working in their primary language.

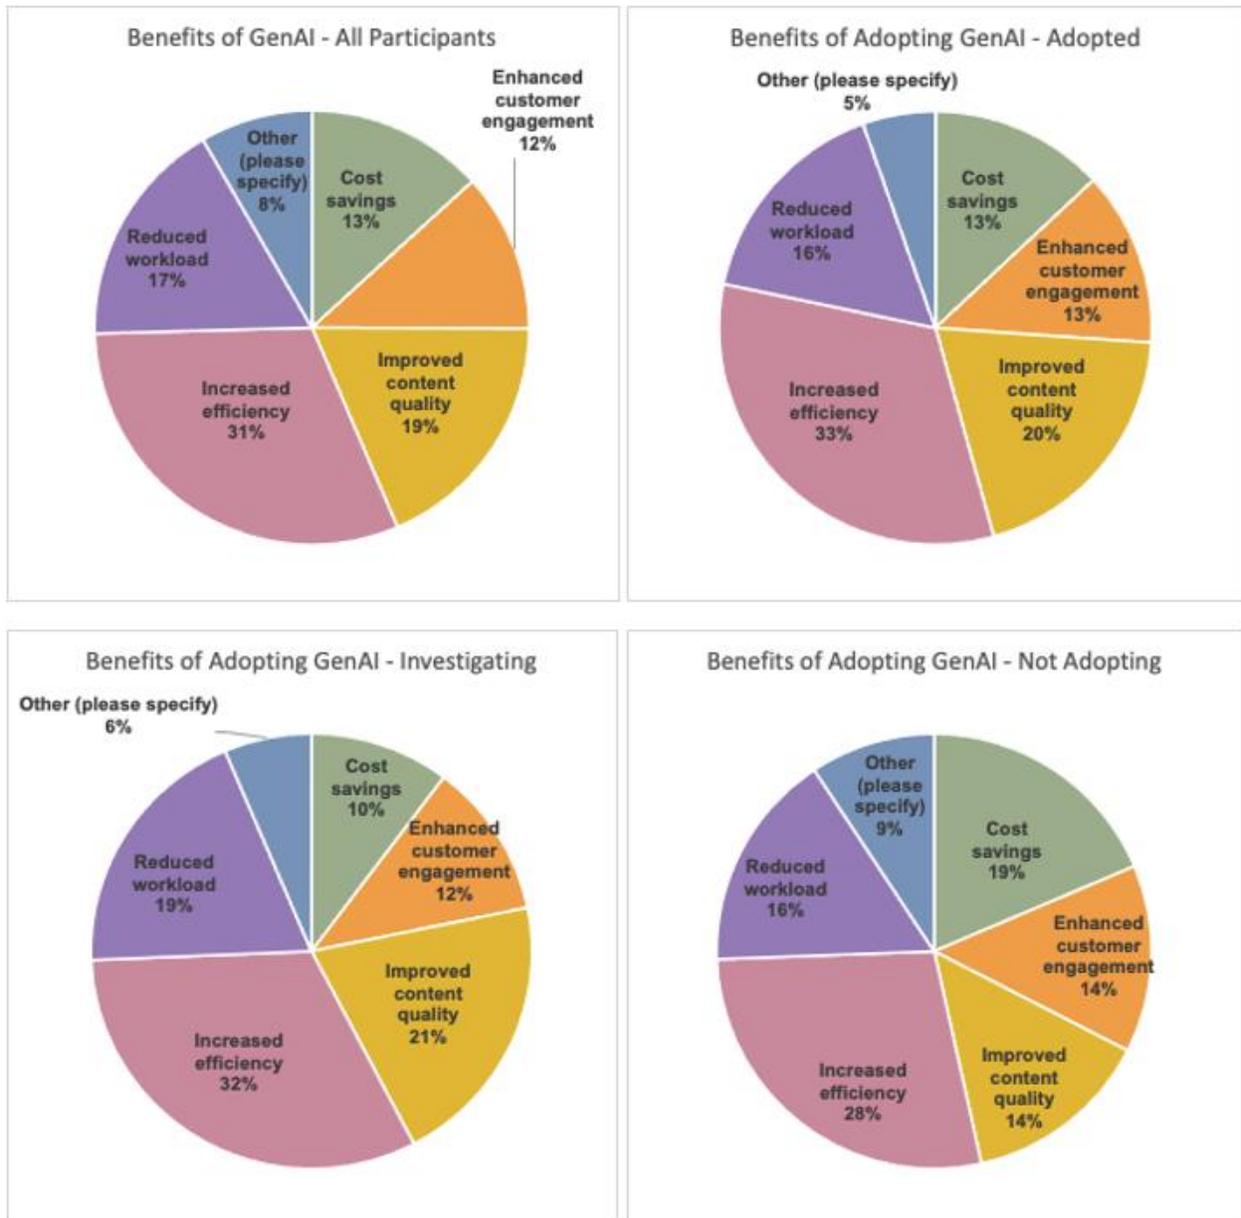


Figure 12: Benefits of adopting GenAI

Promising Applications for GenAI

The promising applications question of the survey looked at how participants were or were planning to achieve the benefits that they are seeking. This question also helps to shape and define ‘increased efficiency’ and ‘improved content quality’ from the benefits question by showing the applications people are implementing.

The most popular application, ironically, was not strict content generation, but rather content analysis and optimization (25% of all participants). This application highlights two of the common sentiments we heard anecdotally throughout the technical communications industry which was that GenAI performs well at summarizing and editing content, but that its ability to generate truly new content was not sophisticated.

Interviewees stated that content analysis included everything from summarizing interviews or information dumps from subject matter experts (SMEs) to consolidating and summarizing request tickets. In both cases, they are using GenAI to change the form and consolidate content that was originally created by humans; rather than to author new content from scratch.

The second highest ranked benefit was automated translation (20% of all participants). Translation tends to be a high and reoccurring cost that many teams are looking to dramatically reduce or eliminate while meeting minimum regulatory requirements and customer satisfaction. There are a few concerns about this as a benefit.

The first concern was that many participants expressed concerns about the accuracy of AI generated content in English (as their primary working language). As the languages become less and less rich on the internet, the ability for GenAI to function successfully goes down. Meaning that using AI to translate languages that are highly prevalent on the internet have a higher chance of being more accurate. However, if the target language was not a rich internet language, then the translation was likely to be less accurate.

The second concern was the need for a translation specialist to be on staff or contracted to review the translation for accuracy. If participants are concerned about the quality of AI generated content in the team's primary language, without appropriate translation support, these concerns should be compounded. It should be noted that, the scenario was not clarified if the AI generation for translation was used by translation service providers. If this was the case, then the translation quality and accuracy risks are mitigated. With that said, none of the participants in our survey identified translation as their industry.

Customer support chatbots are the third most popular promising applications for GenAI (18% of all participants). The anticipation was that these chatbots will be able to accurately and effectively respond to support inquiries automatically. This would drive down call volume and time on calls. The use of automated chatbots (those that can answer a few preprogrammed questions with canned answers or links to documentation) has been a noteworthy assistant for support teams. Particularly when those automated chatbots have been connected to a single-sign-on (SSO) or other audience differentiation method to allow different audiences to receive specialized content.

GenAI powered chatbots are different from their automated predecessors. GenAI powered chatbots, do not use pre-programmed answers, rather they generate new content at the prompt of the user. This means that the answers no longer have to be planned, preprogrammed, and scripted. However, it also means that the supporting documentation is usable by the GenAI chatbot. This includes developing and planning for things like access to the appropriate documentation, as well as consistent and defined structure within the documentation.

What was particularly interesting about the chatbot application was that it moves generative AI towards content delivery. This implies that the content being created will need to model not just

the types of desired answers but will also have to model the style in which those answers should be provided to the end user. The LLM reads not only the semantic tagging of the content but also the tokens in which that content was arranged.

Predictive analytics for content strategy, with 16% of all participants, was ranked the fourth most promising application. It is surprising that this analysis is ranked fourth, when content analysis and optimization is ranked first. Optimizing content is one output of content analysis but so is updating the content strategy. One explanation for this differentiation is that the teams and companies represented in the survey do not all have content strategies in place. Another explanation is that this kind of analysis is the deeper critical thinking that GenAI is notoriously poor at. However, it should be noted that the difference between generating analytics and creating strategy based on the analysis of those analytics are two different things. Analytics are the measures of metrics; they are data, trends, counts, and the like. Analysis is reading the analytics, knowing what they mean, and determining what actions need to be taken to alter the analytics.

General content generation was the last of the main categories to rank (14% of all participants). Given that generative AI is specifically designed to quickly and effortlessly generate new content, it seems counterintuitive that this was the lowest ranking category. The ranking likely reflected the lack of trust that many industry professionals have expressed based on the inaccuracy of AI generated content. Secondly, the low ranking reflects the industry understanding that many in technical documentation are creating content for products that have never existed. This is a limitation of GenAI because if the product has never existed, the predictive nature of GenAI will not effectively generate new content without an accurate pattern for the GenAI to base its generation on. The lower ranking potentially reflected the technical communication industry fear of widespread job loss, particularly if the main application of the AI is to generate technical documentation content.

Several participants included Other in their selection; all were suggested by participants who self-identified as having adopted GenAI or those investigating GenAI. The largest portion of those responses indicated that there were no promising applications in the technical communication industry. Other participants included coding and programming applications including code explanations and code suggestions; such as mathml. The last group of responses moved from text generation to the possibility of image and video generation. Text to image or video generation is a particularly interesting future application of GenAI in technical communications because graphics inclusion can be an expensive and time-consuming part of the content generation process. Using GenAI to assist with animating CAD drawings for teams that are unable to hire full time technical illustrators might be a viable option.

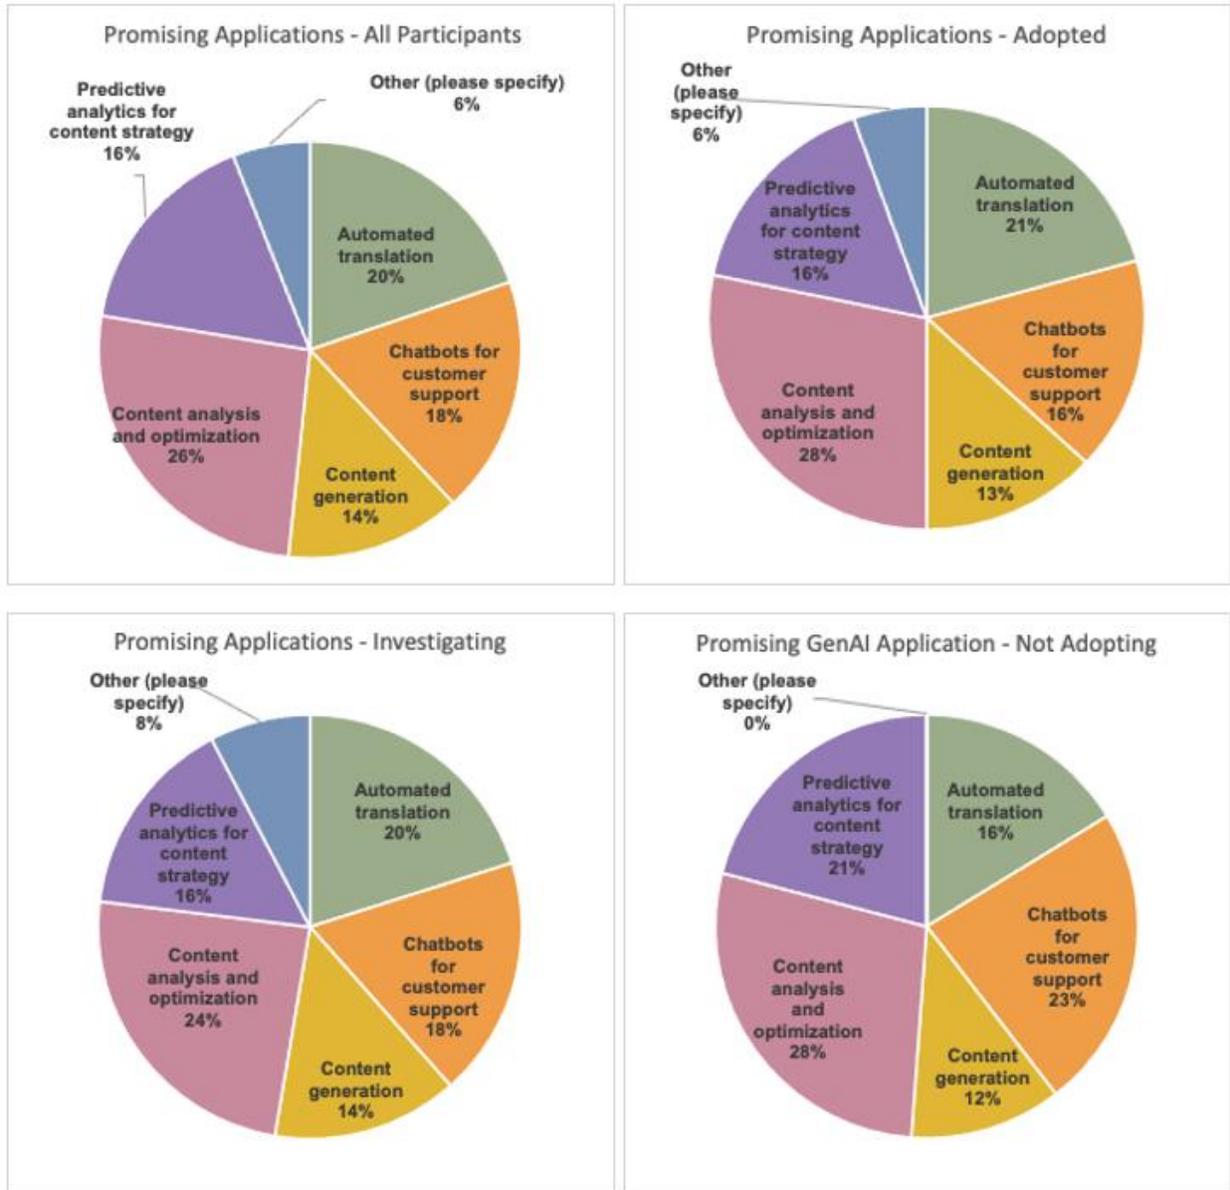


Figure 13: Promising applications of GenAI in technical communications

Content types that will benefit the most

Technical documentation, user support content, and marketing content, across the participant groups, jockeyed for first and second position. Overall, user support content was the content type that technical communication professionals believe will benefit the most from GenAI (22%). Technical communication and marketing content came in tied for a close second among all participants at 20%. (See Fig. 14)

Participants who have already adopted GenAI, however, have a noteworthy difference. Within this demographic, technical documentation was ranked first (26%), user support was ranked second (20%), and marketing content ranked third (19%). The participants in the other two demographics (investigating and not adopting) both ranked technical documentation second and user support first. Marketing content varies more between the demographics. Marketing content ranks third among

early adopters, second among those investigating, and fifth among participants who are not adopting.

Between participants who have adopted GenAI and participants who are investigating GenAI, there exists the same 6% difference in ranking between technical documentation and user support content. We theorize that when exploring GenAI applications, user support is a high initial objective, but by the time implementation is complete, the technical documentation teams are seeing notable benefits for themselves. Potential reasons for this change could be (a) the complexity of successfully implementing a GenAI powered chatbot, (b) the multiple uses of GenAI within the current content creation process, and (c) technical content is more likely to be authored using structured authoring practices which makes this content more GenAI ready. It also means that the true value of GenAI is not being understood in the investigation phase.

Marketing content follows a similar pattern to user support content when examining the three demographic groups. Participants who have adopted GenAI rank marketing content significantly (19%) lower than technical documentation. Those who are investigating GenAI rank marketing content (23%) on par with user support content (22%). While those who are not planning to adopt GenAI rank marketing content as the fifth and lowest ranking (13%). We suspect that some of the same reasons that support content changes rankings applies here. The other suspected reason is that marketing content is further removed from technical documentation than user support is, which could speak to a lack of understanding for how marketing professionals could be using GenAI.

Training materials ranked in the fifth and lowest place for all demographics (including “all participants”), except in the group that is not planning to adopt GenAI. Contributing factors to this ranking include (a) training materials not being authored by the technical communications teams, (b) training materials are generally authored with one or more of the other content types as references pieces, and (c) training materials routinely include interactive content that GenAI cannot create.

The Other category captured some interesting insights into the benefits of GenAI. There were several participants who, again, indicated that there was no benefit to GenAI, but other participants indicated that legal, regulatory, and policy documentation might be a good fit for GenAI technology. Participants who commented further indicated that the repetitive and standardized nature of legal, regulatory, and policy documentation makes it a good fit for generation by AI. Still others called out “internal documentation” as a category. However, without further explanation, it cannot be determined if this “internal documentation” is repetitive of the listed categories because the survey did not distinguish between internal and external facing documents in any of the content types.

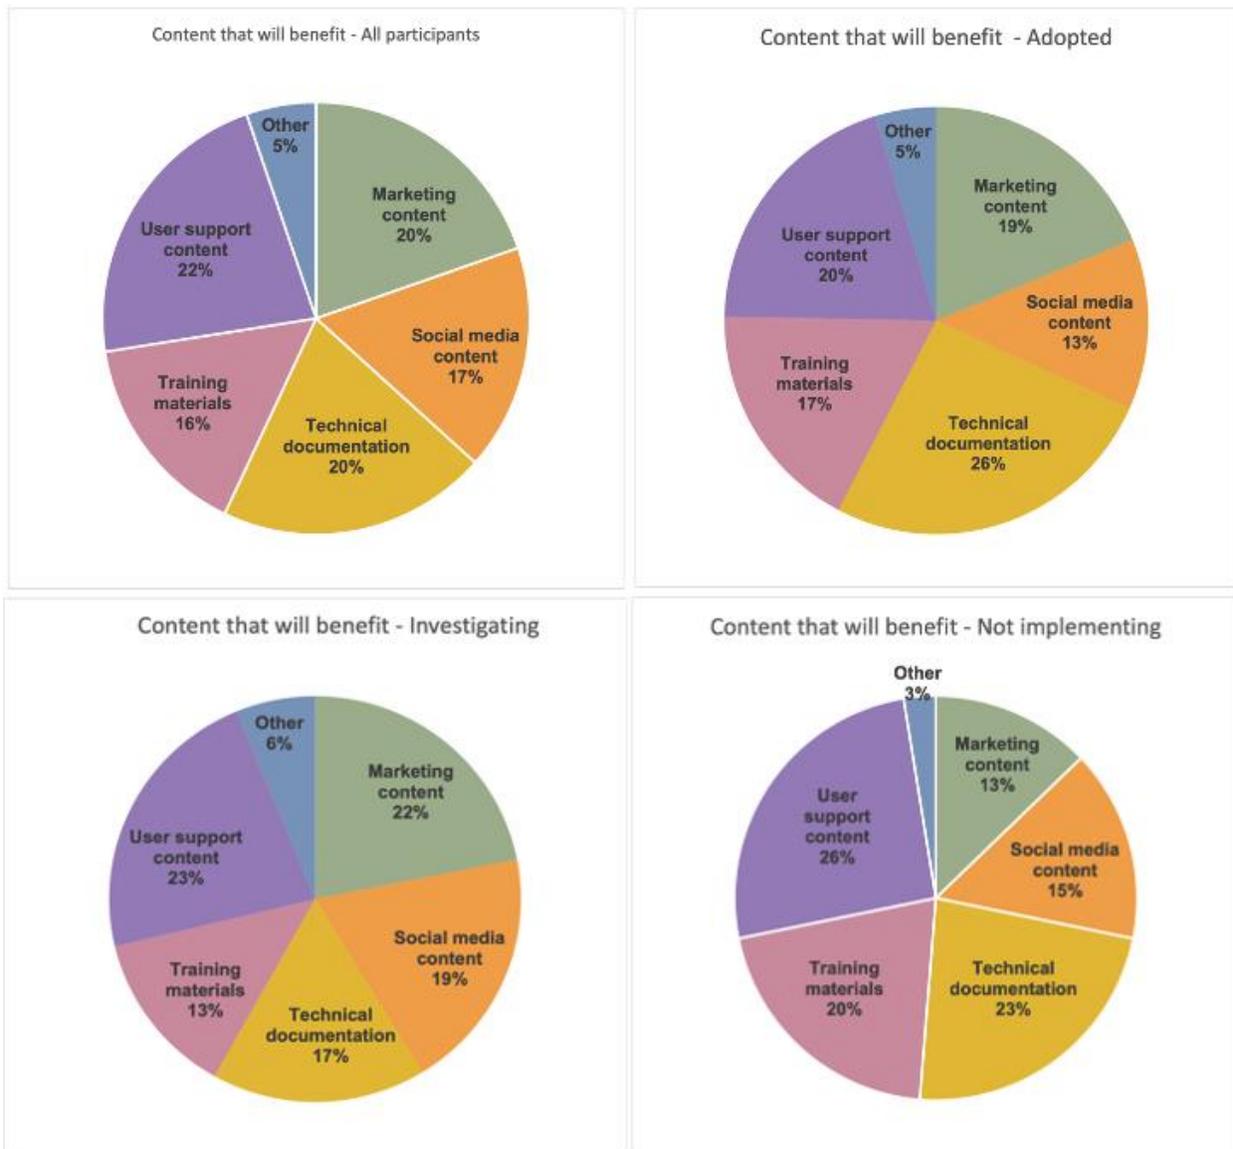


Figure 14: Content that will benefit from GenAI

Challenges to implementing AI

When exploring the challenges to implementing GenAI technology, we limited this question to the demographic of participants that have already implemented GenAI. This was done to eliminate speculation and to focus on real experience. As such, this question was asked to approximately a third of all participants.

Figure 15, "Challenges to Implementation – Adopted," is a bar graph indicating how many of the participants identified each category as a challenge they faced to implementing GenAI. Each participant was able to select as many challenges as applied, including the Other category. This means that the total number of responses to this question is greater than the total number of participants in this demographic and the length of each bar represents the number of participants who identified the category as a challenge. There is no indication of severity of each challenge to

rank if one challenge is more difficult to overcome than another. Participants were able to select all challenges that applied.

The most common challenge was data integration and complexity (14) and the second was lack of skilled personnel (12). These two challenges, if presented in the same project, are going to compound the difficulties of implementation. With 63% of the participants who adopted GenAI implementing a private LLM, the challenge of data integration and complexity speaks to the complications of building private LLMs. It also highlights the importance of clean and well-structured data. Particularly for the technical communication industry, the largest portion of data is the technical documentation content.

Generative AI functions because of patterns and the recognition of those patterns in language and in content structure (semantic labeling). Content with easily recognizable patterns will make AI training smoother and therefore make the implementation of the GenAI easier. Clean and well-structured documentation content means that structure authoring best practices are strictly and consistently adhered to. If the content does not strictly adhere to best practices, consistency in how the content is formed is the second-best option.

Having a lack of skilled personnel in this situation will compound the challenge because it implies that there is no one on staff who will be able to guide decisions and troubleshooting during implementation. Without these skilled personnel on staff, organizations need to decide if they want to push forward with their current staff and resources, allowing staff to develop the required skills. This approach does not come without risks and an extended timeline (which has a cost in and of itself). The other option is to bring in a consultant or partner. Bringing in third-party assistance will likely shorten the timeline and produce more desirable outcomes but does come with an additional upfront cost.

Technical compatibility (9) looks at how all of the parts of the system fit together and how flexible the current state of an organization's software tech stack is. Bringing in GenAI, particularly a private LLM, can include purchasing and implementing several large and expensive tools, such as RAG, the base LLM, chatbot, and so on. Data will need to flow easily and efficiently between all systems. Choosing the appropriate configuration for an organization's tech stack will be critical to successful implementation. Considerations here include:

- Off the shelf versus inhouse built systems
- Cleanliness of in-house created code or connections
- Amount of configuration or specialization that exists
- Current operating versions
- On-premises versus cloud instances

Many of these considerations begin to transition into the next challenge, security restrictions (9). As decisions for the tech stack are made, security of data and intellectual property need to be raised. Of particular concern is the amount of access that public LLMs have to an organization's data and intellectual property. Even organizations that are upgrading tools to access GenAI features are raising concerns about the security of their data when the system relies on a tool that was developed by scraping the internet for content. Many interviewees indicated that their organizations entered into legal agreements explicitly stating that the LLM company will not train their LLM on the organization's intellectual property.

Additionally, many participants and interviewees indicate that they are either not allowed to use the public LLMs or that use of public LLMs is limited to work that does not involve intellectual property. All of the public LLMs, whether they are being paid for or not, will use any and all of the users interactions with the system to further train their product. They are constantly collecting data on how people write, say, and phrase queries as well as whether or not you thought the generation was accurate. This means that while your data is not explicitly ‘saved’ it has been added to a database or ingested into a pattern recognition algorithm. All use of public LLMs is a potential risk to any proprietary content and intellectual property, which makes the security restrictions a notable challenge.

Why these are not the highest challenge says a lot more about the industries and companies that were able to be early adopters of the technology than it does about the severity of the risk. Industries that are heavily regulated and that deal in highly personal data have been slow to adopt; several of them were seen in the “no plans to adopt GenAI” demographic. Companies that work with intellectual property, but not necessarily highly personal or sensitive data, were able to take the risk because the nature of their industry was a mitigating factor. Said another way, the participants who were able to answer this question were able to implement GenAI early which might mean they are biased against security restrictions as a challenge. Organizations who have strong security restrictions were not at a point where they were able to implement GenAI at the time of the survey.

As seen in analysis of the tech stack, the next challenge is high implementation costs (8). Much like security restrictions, this challenge is potentially more severe than its number of indications makes it seem. If survey participants were unable to move forward with their GenAI implementation, they would not have received this question because they would be in either the investigating GenAI or no plans to implement GenAI demographics. It should be noted that even among the participants who have adopted GenAI, cost is still a challenge. In interviews, we saw that cost challenges meant:

- Timelines were slowed.
- Projects were rescope and scaled down.
- Additional phases were added.
- Start times were pushed out for the next funding cycle.

Scalability concerns came next with 7 indications. We infer that while some teams have started to think about scaling the use of this technology, there may not be a clear direction given the newness of the technology and that enterprise level efforts were just beginning. Many of the interviewees were working on GenAI implementations solely for the technical communications team and without the assistance of other teams such as IT, procurement, legal, support, knowledge management, or marketing. This siloed approach does not easily lend itself to scalability thinking because, if scoped well, the scalability for a single team would be something akin to adding additional projects or products to the system. This becomes a change management and system updating task, rather than enterprise scaling of the system to include different types of content.

Encouragingly, leadership for these organizations was very low on the challenge scale:

- GenAI mandated by leadership (3)
- Buy-in or resistance from leadership (4)
- Difficulty with a business case (1)

This is another important potential distinction between organizations that were able to adopt and implement GenAI early and those who have not. In addition to fewer security restrictions, one of the reasons that these teams were able to implement GenAI early was that they had leadership support for the efforts. This implies that there is a potential bias in this data set. It also means that GenAI implementation that have leadership support are more likely to be successful.

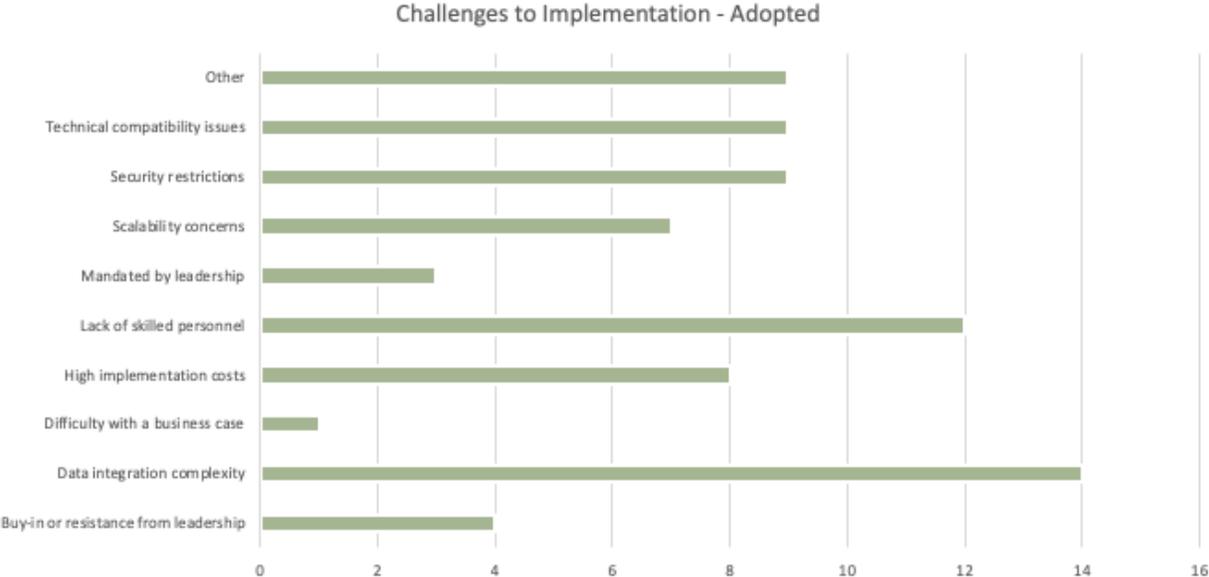


Figure 15: Challenges to implementing GenAI

Why Not GenAI

Reasons for not implementing

While there are many reasons to adopt GenAI, there are also a number of reasons to be cautious about AI implementations. Of our survey participants, 14% indicated there were no plans to adopt GenAI in the next 2 years.

Organizations that were hesitant to adopt, particularly being part of early adoption, were most concerned about the accuracy of content (21%). Tied for second were data privacy and cyber security concerns (17%) along with ethical and legal concerns (17%). Examining these concerns together, we see that these mirror the challenges to implementing GenAI. Additionally, they are the driving forces behind the implementation of private LLM systems and the rising cost of this kind of technology. Even though it is lower on the list of concerns, the lack of skilled personnel (10%) is a compounding factor to the challenges and expense of implementation.

At 12%, the lack of a business case points to three possible reasons: a lack of vision and strategy, industry regulations, or concerns of using the public options. Moreover, these three reasons could be compounding each other in a cycle of cause and effect. For example, if regulated industries were not strategizing toward GenAI, they may not have the skilled personnel to start implementing the project. This, in turn, means that the business cases could be more difficult if it needs to include additional personnel or is simply too large of a project for the size of the team.

Of particular interest in the Other category was the comment that the participant’s content is not “AI ready”. This point is important to discuss. Organizations that are already working in structured authoring with an implemented CCMS and matured processes are more prepared for a successful implementation of private LLM for a sophisticated GenAI implementation. If an organization is not working in structured authoring, the best practice would be to make that transition first.

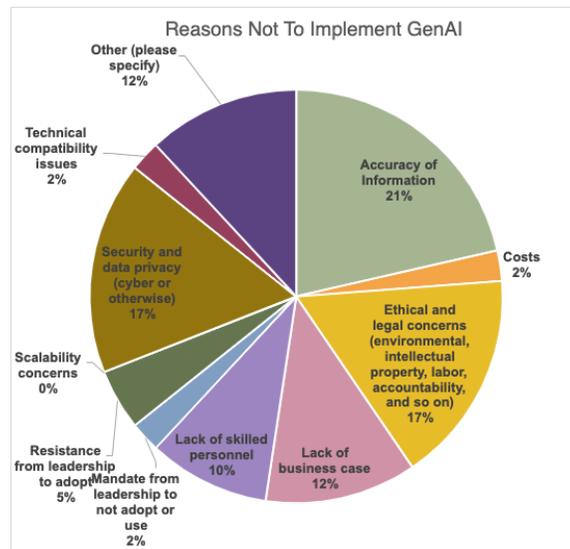


Figure 16: Reasons to not implement GenAI – Not Implementing

What would change the decision

Following the reasons that organizations are not exploring GenAI is the question: “What would need to happen to change the decision to not investigate GenAI?” We wanted to know what, if anything, it would take to change the minds of decision-makers when it came to GenAI. We were hoping to understand the root cause(s) of not implementing GenAI within the next 2 years.

The biggest factor to changing the decision is improved technology (22%). We did not get so granular as to understand which technological improvements participants were looking for the most. However, we infer that the improvements should focus around content accuracy, reliability, and security as those are the common challenges to implementing and the top reasons for organizations to not implement GenAI.

In a similar vein, mitigation of ethical concerns (91%) came in as the second driving factor to changing the minds of decision-makers. Participants expressed notable concerns about GenAI such as data privacy, cyber security, and accuracy of information. See [Ethical Considerations](#) for further explanation. The top ethical concern is that using the public LLMs means that the user agrees to allow the LLM to train on any content and interaction between user and LLM. This concern is driving the move to private LLMs. Private LLMs are based on the public ones and therefore there are additional ethical concerns about backdoor access to the client content as well as the legality of contractual agreements. A second ethical concern is that the public LLMs were created and trained by scraping the internet for content; the legality of which is still being debated in the US justice system but has also spurred the art community to establish programs and practices that prevent GenAI from successfully using their work.

The third change for the decision to be reconsidered is customer demand (16%). Customer demand, as part of audience analysis, has always been a driving force in the technical communication industry. For GenAI, this can be a large ask and heavy demand on a technical communications team and organization because the scale and cost of implementation can vary greatly. The scope and application of GenAI varies with the scale and cost, meaning that although customers might want GenAI and might even become accustomed to using it does not mean that it is a good technical communication solution. A good solution is one that addresses the root concern of, in this case, the audiences; which is not by default a lack of GenAI.

Industry trends (6%) and competitors are doing it (9%) come in rather low in the survey. Similarly to audience demand, just because it is a trend or because your competitors are doing it, are not likely to provide enough meaningful information for your leadership to change direction. They are points in an organizations timeline to pause and reevaluate the decision to not bring in GenAI, but how much time is spent on a cost benefit analysis, as well as root cause and alternative solutions analysis have many other business dependencies.

Several participants responded to the Other category. As previously discussed, a few of the answers alluded to improvements in the technology and privacy of information. In addition to those comments, other participants reflected on the need for more personnel. Given that adding personnel to the team is not a common goal of GenAI implementations, it was not given a place on this multiple-choice list. There are two inferences that can be made from this answer. One option is participants are pointing to the need for personnel trained and experienced in GenAI to make the implementation and maintenance of such a project more manageable. A second option is that if the technical communications team required additional people to keep up with the current and future workload, that GenAI might be worth investing in.

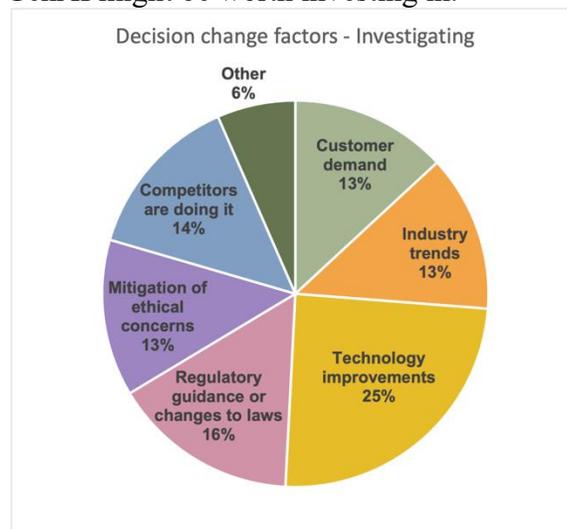


Figure 17: Decision change factors - Investigating

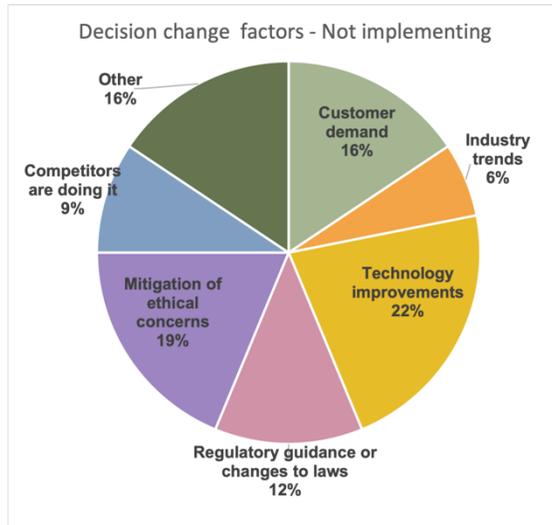


Figure 18: Decision change factors - Not

GenAI and Content Creation

How has it changed the content creation process

For all the hype of GenAI and its supposed ability to revolutionize content creation, we did not see such dramatic results. While there have been improvements in the content creation process because of GenAI, we would not say that it is revolutionizing the industry. With that said, we are excited by the moderate improvements that 52% of participants who have adopted GenAI have experienced. (See Fig. 19.)

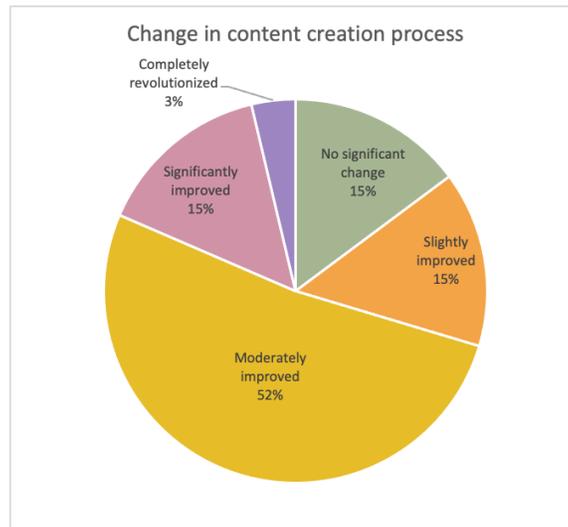


Figure 19: Change in content creation process - Adopted

The follow-up interviews were particularly helpful in understanding the differences in the degrees of improvement. First, it should be noted that there were no direct correlations between any specific model or implementation type and the degree of improvement. While some participants shared what model they were using, not everyone did and there were several different models used including private LLMs. Secondly, all of the implementations were new enough that no one had

clean and reliable data to track specific parts of the implementations to specific gains made in content creation.

What was common throughout the interviews is that as a group, all teams that have implemented GenAI have experienced at least some improvements to all steps of the content creation process. A majority of the improvements were seen in the early stages of content creation. In particular, the stages including gathering subject matter expert (SME) information and initial draft creation. The teams we spoke to indicated that summarization functions as well as formatting functions were useful. On the other end of the content creation process, at publish, there were also noteworthy gains made by including GenAI chatbots. The GenAI chatbots are driving down the costs of support by providing answers quicker, and simultaneously, increasing content quality.

Teams that implemented a full-scale GenAI solution with GenAI as a writing assistant and as a content delivery method were able to create a circular process for content improvement. Technical communication was able to use GenAI to improve content quality, push content to users via a GenAI chatbot, and then use errors found in the chatbot to improve the corpus of product documentation when answers were not found or when the chatbot provided incorrect answers. Either way, errors in product documentation were identified and corrected to improve documentation quality. Additionally, support groups supported by a better chat bot saw improvement in quality metrics such as call reduction and decreased time per call.

Furthermore, with all of these gains, not one team sacrificed quality or removed steps of the content creation process. It seems that the sum of the parts is greater than the whole. GenAI was not used to replace any essential function of a content creation team; it was used as a tool to make the people in those functions more productive. In making them more productive, it opens their bandwidth to the deeper knowledge management and creation tasks such as information architecture, audience analysis, customer journeys, and other process maturity tasks. (See Fig. 20)

Adding GenAI to the content creation process is impacting content strategy. How and where the content strategy changes depends on the implementation. Upgrading tools to include GenAI are having minor impacts to the strategy. However, implementing GenAI as a content delivery mechanism has a larger impact as it requires a more robust plan to implement, monitor, maintain, and verify.

Post-GenAI Content Creation Process



Figure 20: Content Creation Process with GenAI

The improvements to the content creation process are slow to impact scalability of content creation. While the initial drafts creation times are down and overall time to market is getting faster, it is not at rates that are allowing teams to take on larger scales of work. We saw large numbers indicating moderate improvement (10 participants) and slight improvement (7 participants) to scalability from teams that have adopted GenAI. However, approximately the same number of participants (8) have seen no impact on scalability. (See Fig. 21)

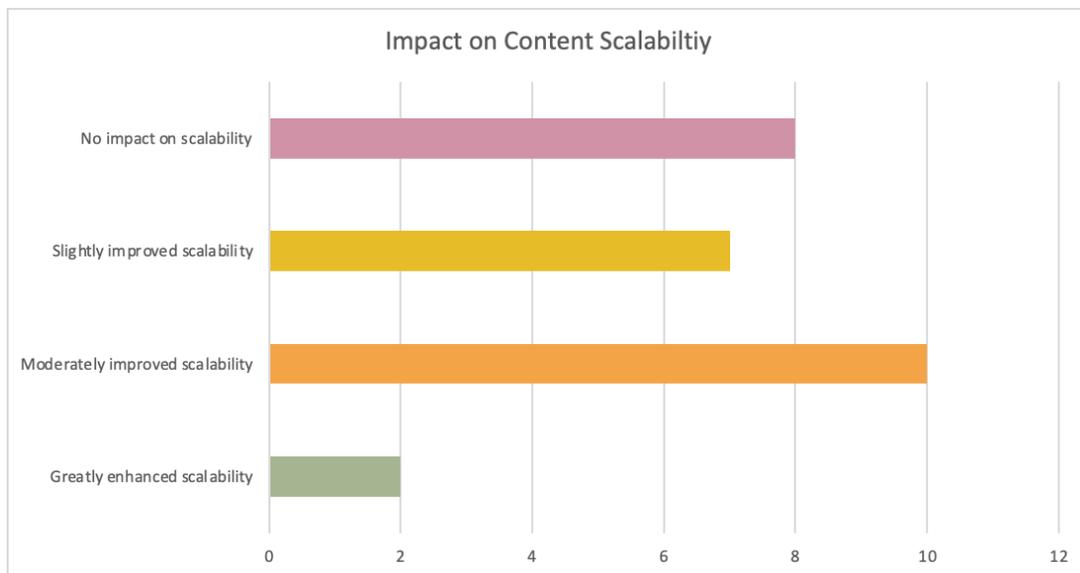


Figure 21: Scalability of content - Adopted

This limitation on scalability of content stems from three reasons. Primarily, as we discussed, the content creation process has not changed drastically and maintains the same number of humans in the process. This means that the people required to complete the process has remained constant as have the steps in the process. With only incremental improvements seen in the process, it is unsurprising that there are limited impacts to scalability.

Second, many technical communication teams are under-staffed and struggling to complete their current workload. Implementing GenAI on these teams has likely "slightly impacted scalability" because it means they might have the ability to successfully maintain the status quo, but not necessarily take on additional responsibilities. A common refrain that we heard in the interviews is that with the capacity opened up by the implementation of GenAI, teams can address tasks on their backlog; again pointing to the fact that they are struggling to meet their day-to-day workload, let alone have the ability to expand their capacity solely with the implementation of GenAI.

Third, in large scale and full implementations where we are likely to see the greatest opportunity for scalability, we are also seeing the dip in productivity that is part of large-scale software implementations. This dip in productivity is due to the learning curve (time spent learning how to best use a new tool) and the division of time and labor between implementation activities and those required to maintain current workload. In the years that follow the initial implementation, ROI in time, money, or resources would expect to be seen; this would be an impact on scalability. However, many of the folks that implemented GenAI, have only been in their first year of implementation and therefore would not expect to be seeing the scalability benefits.

Factors for why it hasn't changed the process

The reasons that there have not been significant changes to any one step in the content creation process or the elimination of steps from the process reflect the challenges that many teams face at implementation. Lack of skilled personnel, privacy concerns, data challenges (security and quality), and information accuracy in addition to ethical concerns all mean that throughout the content creation process a human will remain in the loop. Keeping the human in the loop allows for verification at each step of the content creation process. Verifications ensure that quality is established and maintained throughout the process. Humans in the loop also maintain traceability, accountability, and explainability by quickly identifying changes in quality.

Given that the teams we interviewed were early in their implementations, we also heard that content was taking longer to produce with GenAI than it was to produce it manually. Learning curve, implementation, and full integration into the content process are leading reasons why GenAI is hindering the content creation process in some cases. Similarly, we are also seeing where additional steps are required to verify the quality of content, which is adding steps to the process, even if the total time to market is decreasing.

In most instances, teams are using GenAI as a tool to close a learning gap and allow the content creation team to work faster, rather than changing the foundational content creation process. For example, folks who are not full-time content creators can get a jump start on creating new content when they need to which will give the content creators a better initial draft to work with. This includes being able to do better and faster research as well as adding structure and formatting to unstructured content.

The factors limiting process change are also impacting collaboration. Some teams are seeing improvements in their ability to collaborate, but when examining these improvements against the small changes to the stages of content development, we infer that improvements to collaboration are likely seen in the improvements of the initial draft. By lowering the bar of initial draft creation, teams can work on deeper knowledge work or have more productive conversations about the accuracy of content, rather than being distracted by the formatting, structure, or organization. It

also demonstrates that the learning curve to get to initial draft is lowering because people can do product or style research on their own.

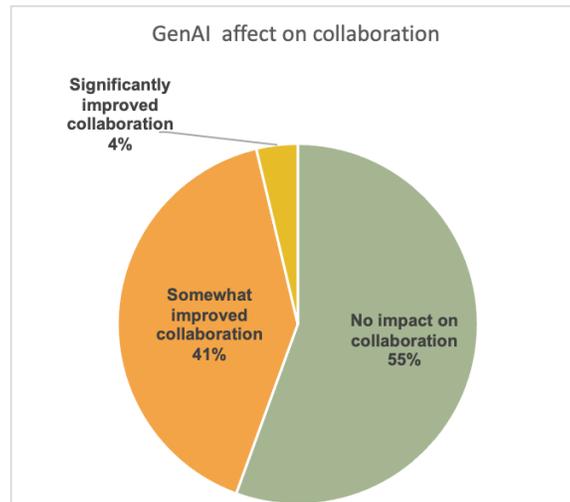


Figure 22: GenAI effects on collaboration - Adopted

Evaluating GenAI

As a tool

There are many metrics that can be used to measure and evaluate GenAI. Each organization needs to determine their reasons for implementing GenAI, and then using those reasons, determine what metrics are the most significant to the success of their project. Knowing this, we did not ask about specific metrics in the survey. Rather, we asked broadly how teams were validating the GenAI implementation.

During the follow up interviews, none of the interviewees were tracking individual or specific metrics. Instead, they were relying on whether it was “working” and if the content was “accurate.” When discussing their quality assurance processes, answers aligned to the broad, holistic, metrics that we saw in the survey. Some interviewees discussed that they had defined “accuracy” in their content strategy as part of the manual or pre-GenAI process and that any content generated by AI must meet the current minimum quality standards.

One thing to highlight in this, however, is that content strategies have not likely been updated for verifying if the GenAI system, especially with RAG, is functioning correctly. With early adopters going all in on an GenAI solution quickly, it did not appear that they established what the criteria or requirements for the GenAI was other than safe, reliable, and accurate content. This ignores metrics of LLM and RAG functionality such as accuracy, precision, or efficiency.

The most plausible explanation for this miss in metrics is that these are metrics that can and should be measured by the team that is installing the LLM. In many cases this is the responsibility of an information technology (IT) team. While technical communicators are leading and championing efforts, like all enterprise level software implementations, IT plays a strong role in the effort and

is responsible for choosing an appropriate base public LLM. To do so, they would have needed some metrics around accuracy, precision, efficiency, and other upon which to base their decision. If this is indeed the case, then those metrics likely exist and are being monitored, but perhaps not being communicated to the technical communications department. Similarly, if IT is focused on the system aspect of GenAI then it makes sense that the technical communication team is focused on the quality of output, which would need to match the quality standards set forth in the content strategy because the quality of the content will not depend on the way it is created. In sum, this explains why 10% of technical communication teams have no formal validation process. (See Fig. 23)

The survey reveals that 45% of participants are relying on manual reviews by experts. This could be the SME, a professional editor, or similarly qualified person. Essentially, this is one of the main roles of the human in the loop – to verify the type and quality of content throughout the content creation process. Secondly, we saw that teams were cross-referencing outputs with other reliable sources (33%). In interviews and throughout the survey, we saw a number of participants using GenAI as a research tool. Along with verifying accuracy of a specific piece of content against other sources, verifying the research output of GenAI is another reason that cross-referencing comes up as high as it does. (See Fig. 23)

The 10% listed in the other category for GenAI validation, predominately reflect the previously listed categories of manual review by an expert and cross-referencing content with reliable sources. Non-specified “testing” and “feedback from users” were also listed. One participant indicated that they are “developing knowledge graph driven verification.” Of the responses we received, this would be one of the more elegant verifications.

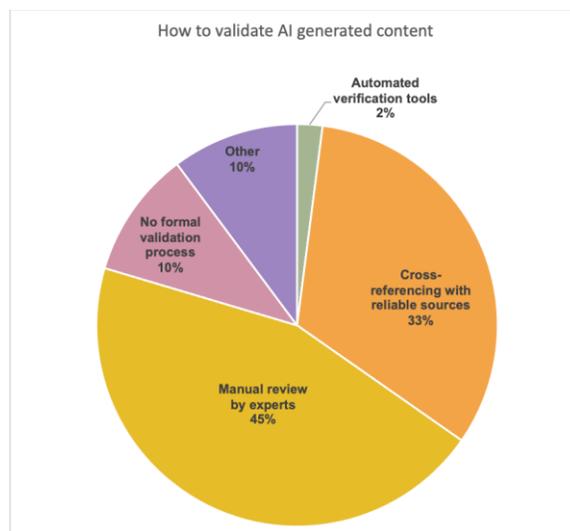


Figure 23: Validating GenAI - Adopted

Another critical measure of successful implementation is employee feedback. This can look like an online form, meetings, any number of open forums, or simply that they are willing and able to use the tools. Teams that have implemented GenAI are relying on the feedback from employees using the new tools and functions to understand what is working, how well it is working, and what should be improved. Teams that are using writing assistants are relying on skilled writers to know if the tool is adhering to the style guidelines. Teams that have released a GenAI powered chatbot

are collaborating with cross-functional teams such as support to understand not only the writing style that the chatbot is using but also where the gaps in the product documentation exist.

Perhaps the most important and strongest measure of the GenAI implementation is pilot testing. In all cases of pilot testing, technical communication professionals devised scenarios and use cases that tested the quality concerns they had with the GenAI system. For writing assistant tools these scenarios focused heavily on the challenges the team had in writing content or some of the more esoteric requirements of their style guide.

For teams that implemented a chatbot, many of them tested the chatbot with internal teams such as support or customer service. This allowed the support team to query the chatbot with real life customer inquiries to verify that the chatbot provided the correct answer and did so in a manner that was appropriately communicated. By partnering with support teams, use cases were written based on the most common reasons support was called or common searches run through a help site. By using real cases, technical communication teams were able to improve the corpus of their documentation and help support meet their objectives such as call deflection and decreasing time on calls by providing more or better documentation on the high-volume call use cases.

Lastly, teams were demonstrating success of the implementation by highlighting increased capacity of the team. We heard from interviewees that they were able to show items being ticked off the backlog, working on projects that have been on the “nice to have” list, and taking on additional product or project work without having to increase headcount. Moreover, by presenting this as an increased capacity, it helps to demonstrate the value technical communication teams bring to the organization. These become “cost avoidance” arguments rather than “cost saving” arguments during ROI discussions; particularly in the first year or two where implementation costs and learning curve costs are still being incurred by the organization.

Ethical considerations

Another evaluation of GenAI is through an ethical lens. While these will not necessarily help make distinctions between two LLMs, they have impacted decisions on whether or not to implement and what kinds of implementations organizations are willing to invest in. Many of the ethical concerns mirror the challenges to implementing GenAI as well as drivers to change implementation decisions. It is also important to note that these ethical concerns are not a deviation from those traditionally held by the industry.

Across all three groups of participants, information accuracy (3.6 score), privacy and security (3.4 score) and intellectual property (3.3 score) top the lists, while labor (2.6 score) and environmental impacts (2.2 score) fall to the bottom of the list. The top three ethical concerns are the ones that have traditionally been part of the technical communication industry mantra: the right content, to the right people, at the right time. So even though the means of content productions have changed, the primary ethical considerations of producing that content have not.

Similarly, labor and the environment being at the lower end of the concerns is unsurprising because they have not been traditionally considered as part of providing documentation. Product documentation is a requirement for successful use of the product and, in many cases, is demanded by regulations. However, the exponential demands of both labor and ethics with the invention and widespread adoption of GenAI should become a larger consideration for organizations.

Middling in the ethical concerns are transparency (3.2 score), bias (3 score), and accountability (3.1 score). Again, building on tradition, these are all qualities that the technical communication industry has relied on to consistently create high quality content and skilled content creators. Changing the tools that are used should not change the value that the industry places on these ethical considerations. However, transparency is one of the concerns that need to be revisited with GenAI because the technology itself is changing what should be transparent to whom. At stake here is the question of what AI generated content should be identified to the end user and is it important to the end user that they be told what content is produced by GenAI. Moreover, global governments are making regulations that are in the early stages of defining these points. Nonetheless, what content gets indicated, how to indicate, and when to indicate AI generated content should be part of any content strategy that addresses GenAI implementation.

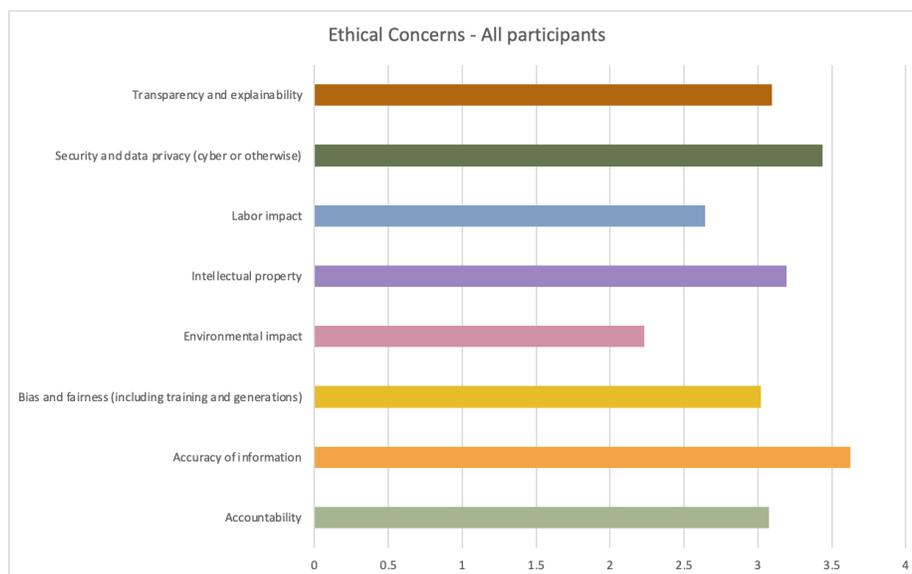


Figure 24: Ethical concerns about GenAI

Trends

There are several promising trends for technical communications and GenAI. First and foremost, is that technical communications are going to be in the forefront of the GenAI implementations. The essential reason for this is that high quality and consistent product documentation is at the heart of all GenAI systems. The second promising trend for technical communications is that our traditional and well-established content creation process will evolve but not wholesale change.

With 68% of our participants indicating that they are leading or championing GenAI implementation efforts, 73% of implementations are private LLMs, and 94% of those private LLM implementations training, technical communication teams are essential for successful GenAI implementations. It is the technical communication professionals that have established the processes to create and maintain the sophisticated product documentation that the GenAI systems require.

Validation that the technical communication content creation process is mature enough to handle this kind of tool change is in the fact that none of the steps have been wholly removed from the process. Each step has benefited from the addition of GenAI, but there remains the human

verification, creativity, and deep knowledge work required to produce content. Looking ahead, the need to produce content for new products, those that have truly never been created before, will also need technical communicators to produce the first documentation set. Then, with that newly created documentation set, the GenAI will need to be maintained and retrained in some cases with the new content. This will assist in avoiding content degradation because it will continually feed the GenAI new content to add to its statistical algorithm.

Conclusions

The technical communications industry will be significantly impacted by GenAI, but this is not a bad thing. Jobs will be intact, but in new and exciting capacities that involve more strategy and deeper knowledge work because many of the rote tasks will be made easier. For example, product documentation requirements will need to change to accommodate GenAI as a user. This means understanding GenAI as a consumer of content and what that means for product documentation as well as content strategy.

Technical communicators are leading the GenAI efforts. There is going to be growing interest in what GenAI can do for organizations and technical communication professionals will need to understand how to create GenAI-ready content, how the systems work, and how to validate their content in a GenAI system. Much of this is not going to be the hands-on keyboard work that many are accustomed to doing. It will, instead, be deeper strategy and information-as-knowledge work required to maintain a complex system and to create documentation for products and equipment that are newly designed.

GenAI is not something that should inherently be embraced or discarded. It is a powerful new technology that requires vision, planning, skill, and leadership to determine if and how it can and should be implemented at an organization and in a content creation process.