

AI Assistant Whitepaper




Table of Contents

1. [Introduction](#)
2. [Objectives](#)
3. [Goals and Vision](#)
4. [Application and Use Cases](#)
5. [Technical and Functional Description](#)
6. [Security and Privacy](#)
7. [Challenges and Future Directions](#)
8. [Conclusion](#)



1. Introduction

The emergence of generative AI has significantly reshaped the technological landscape of 2023. This groundbreaking advancement can be attributed to the early experiments in neural networks conducted in the late 20th century. Since then, generative AI has experienced substantial growth, impacting a wide range of industries, from creative arts to software development.

The increased adoption of generative AI this year has solidified its position as a pivotal component in the technological arsenals of forward-thinking companies. While these tools have unlocked new levels of creativity and productivity, they also present new challenges in terms of ethical considerations and data privacy. To fully harness the potential of generative AI, companies must carefully integrate these tools into their existing frameworks, ensuring they enhance rather than disrupt current operations.

To effectively engage with generative AI technologies, it is essential to design interfaces and experiences that are not only powerful but also intuitive and user-friendly. These advancements empower organizations to leverage creative and operational efficiencies while navigating the complexities of generative technology.

Through the adoption of advanced generative AI solutions, organizations can redefine how creative and operational tasks are approached, leading to more innovative outcomes and greater efficiency. This strategic integration can drive significant competitive advantages, enhancing both productivity and market relevance. Embracing generative AI represents a crucial step for companies aiming to excel in the dynamic digital arena of today and beyond.

However, the use of generative AI comes with its share of risks, including data misuse, biased outputs, and the potential for creating misleading information. To address these concerns, Templafy is focused on delivering managed AI services that prioritize safety and compliance. By creating a controlled environment, Templafy assists organizations in managing the risks associated with generative AI, ensuring the responsible and effective use of these powerful tools.



2. Objectives

In this document, we aim to provide you with a comprehensive overview of the product. Our primary objective is to give you a technical and functional understanding of the Templafy's GenAI, including its underlying model and architecture. We will also outline the main functionalities that end users and administrators can interact with when using the product.

We understand that practical use cases and applications are essential to understanding the value of our technology. Therefore, we will describe some of the ways Templafy's GenAI can be deployed to enhance processes within your organization.

As a security-first organization, we also understand the importance of privacy and security. We will discuss the security and privacy implications of Templafy's GenAI. We want to be transparent with you about the risks that your organization might incur when integrating this product into your tech stack and show you how we have been working to mitigate these risks.

Lastly, we will offer insights into Templafy's vision for what is to come regarding AI. We are confident that our technology will continue to evolve and provide even more value to our customers. We hope that this document will give you a better understanding of our GenAI tools, their capabilities, and how they can benefit your organization.



3. Goals and Vision

Templafy enables knowledge workers to create high-quality and trustworthy documents at speed, by blending rule-based automation with managed AI generated content, giving them and their organization control, confidence, and peace of mind.

GenAI is being used by employees without companies' permission. Companies do not have control over the quality and accuracy of the documents being generated. Scalable use of AI to boost end-user productivity can be expensive and tools might require a lot of expertise.

With Templafy's AI, companies can allow access to time-saving generative AI functionalities within the employee workflow. In this way, companies can have centralized control over when and how AI generation is used within document generation. Managed AI together with rule-based automation embedded into employee workflows increase document efficiency and quality, scaling to specific needs of the employee.

Automated company-approved document standards, data, and content, together with AI-generated content, will help us to eliminate manual document creation and allow employees to focus on the knowledge work they were hired to do.

Security and privacy are arguably the main aspects to consider when evaluating AI solution providers. This document summarizes how Templafy handles those aspects in the context of Templafy's GenAI applications.

Our main GenAI capabilities are the following:

- **Centralized control over AI content output:** Create and customize the AI actions available to users of the AI Assistant
- **Centralized tone of voice:** Define the tone of voice in a central place and use it in all relevant actions, ensuring consistent language that does not go against your brand
- **Connect to custom LLM model setups:** Connect the AI Assistant to a regular or fine-tuned models on Azure OpenAI Service, ensuring control over usage costs, security, and data privacy - Included in Custom AI Assistant
- **Enterprise-grade security and data privacy:** Regardless of setting up a custom AI connection or not, requests from users are handled by the model on Azure OpenAI Service in a limited and secure manner.



4. Applications and Use Cases


The AI Assistant covers a myriad of use-cases due to the ability to create actions using any prompt and connecting to **any model stored on Azure OpenAI Service**. Currently, when first enabling the AI Assistant it already covers four different use-cases. In this section, we will go in depth into what each default action can achieve and what is possible when creating actions that use custom prompts and instructions.

Document compilation and document editing can benefit from Generative AI if managed well


Document compilation

Template automation

For brand integrity and policy adherence
 Primary objective: **Compliance**




Automation type needed: **Rules based**




Rules based automation will be required to ensure correct and up-to-date use of mandatory company information, classification labeling and company branding

Business data automation

For speed and accuracy
 Primary objective: **Data accuracy**




Automation type needed: **Rules based**




Rules based automation will be required to achieve 100% accuracy in business data from data sources. **AI** will help locate data sources and generate supporting content based on the data

Content automation

For optimal productivity
 Primary objective: **Employee productivity**




Automation type needed: **Managed AI**



AI can help unlock the full productivity potential of document automation. Companies will need to **influence the AI through APIs and managed prompting** to ensure high quality and trustworthy outputs

Document editing



Content editing

For efficient customization
 Primary objective: **Document performance productivity**

Automation type needed: **Managed AI**

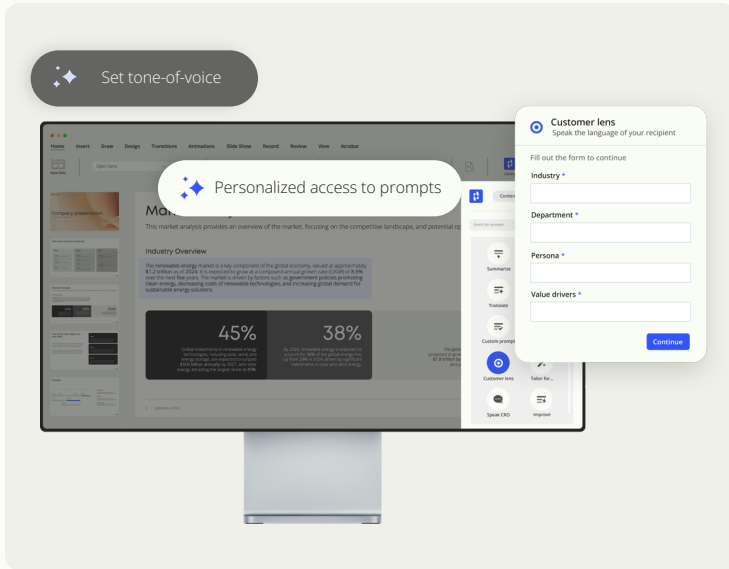


AI can serve as a powerful capability to help with iterations and editing of content. Companies will need to **influence the AI through APIs and managed prompts** to ensure high quality and trustworthy outputs



4. Applications and Use Cases

AI Assistant



Increase employee satisfaction by providing the latest technology



Access to company-approved GenAI tools within existing workflows in PPT, Word, Outlook and Excel



Create and improve content with simple, pre-defined actions as well as custom prompts from your own library



Define and maintain a consistent Tone of Voice and writing style across all AI-generated content

Maximize your AI investment and quality of AI generated output



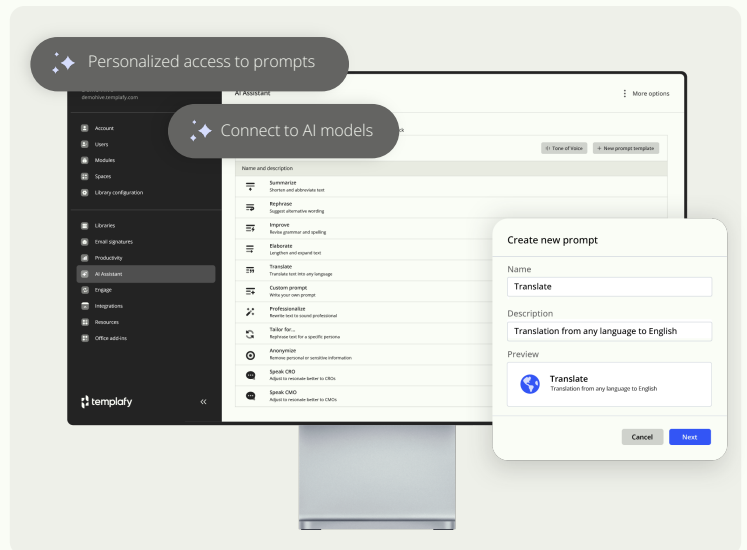
Customize and filter prompt libraries based on users to ensure employees have access to most useful actions tailored for them.



Link your preferred AI models and RAG to prompt action for best possible output quality and increased ROI of your AI investment.



Stay compliant with the highest international enterprise cloud security standards.



To effortlessly begin using the AI Assistant, pre-set prompts/actions are provided along with easy-to-use starter templates. All prompts can be **modified or removed by customer Admin at all times.**



4. Applications and Use Cases

AI Assistant

To effortlessly begin using the AI Assistant, pre-set prompts and actions are provided along with easy-to-use starter templates. All prompts can be modified or removed by the customer admin at any time.



Ready-made prompts

AI Assistant comes with ready-made prompts automatically enabled for the end-users. These prompts are designed to help employees in editing, creating, and adjusting the formatting of their text-based content. Prompts will inherit the tone-of-voice guidance created by your organization and can be edited further through the admin center.



Starter prompt templates

Starter templates are ready-made best practice prompts created by our engineers based on customer feedback and product research. These prompts can be chosen from the Admin Center and modified further to fit your specific needs. The Templafy team will continue to add tested, time-saving prompts to help you get started quickly and make the most out of your solution.



Custom made prompts

Creating custom actions to facilitate any use case important to your company comes with full customization of the action's title, description, instructions/prompts (system message and user message), and the ability to pick or upload an icon for the action. The tenant admin can manually add tone of voice to prompts in the admin center, allowing for flexibility in prompt creation and the option to omit tone of voice from specific custom actions.



4. Applications and Use Cases

Templafy and Copilot for Office 365

Templafy has been working closely with Microsoft for nearly a decade, solving enterprise challenges together. This deep collaboration ensures that Templafy integrates seamlessly with Microsoft's suite, including Copilot for Office 365 – ensuring a smooth user experience and an uninterrupted workflow.

Templafy provides a robust document management layer that enables enterprises to harness the productivity power of Copilot's AI – without sacrificing control of brand consistency and compliance.

Create documents through Copilot for Office 365

Create semi- or fully automated documents with Templafy and Copilot for Office 365. Copilot searches for the best presentations and documents to fit your use case. Templafy compiles documents utilizing company-approved content and information from third-party applications and Templafy together, utilizing Copilot.

Create compliant content with Copilot for M365

Start all documents following company guidelines. Templafy automatically creates documents with company-approved templates, including branded logos, styles, legal disclaimers and metadata based on the user's profile information.

Introduce document automation within Copilot process workflows

Automate document creation in all process workflows with Templafy Document Generation API. It is easy to create a workflow with automated system triggers in the Copilot studio. The Templafy Document Generation engine can populate the correct template with up-to-date data from any application to create either a partially or fully automated document.



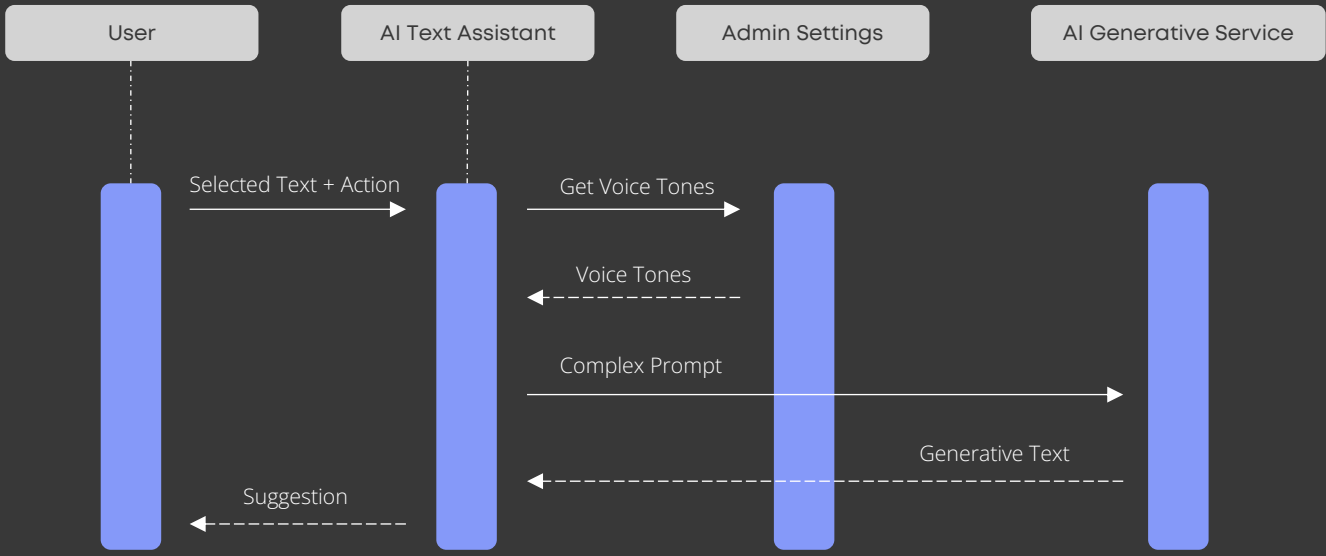
5. Technical and Functional Description

Data Flow Diagram – AI Assistant

The flow starts in the user's Outlook, Word, or PowerPoint. The user selects the text that they want to edit using the AI Assistant by highlighting it within the application.

The user can then open the Templafy add-in and select the AI Assistant tab. In this tab, all the active prompts will be displayed, and the user will be able to select the one they prefer. The AI Assistant will get the voice tones that were previously inputted by the tenant admin in the admin center. A complex prompt containing the selected text, the selected action and the tonality provided in the admin center is then created.

This complex prompt is sent to the AI generative service, which in the current set up is Azure OpenAI Service, which will generate the output text and send it to the AI Assistant. At that point, the output is displayed within the Templafy add-in, and the user can either select it or ask the AI Assistant to generate a new output. Once the text is suitable for the user, they can click on "Use text" and the text will be inserted in the email, document, or presentation instead of the text that had originally been selected by the user.



5. Technical and Functional Description

Model and Service Hosting Locations

This is only applicable if the customer decides to leverage Templafy's model. If the AI Assistant is connected with a custom model, the specifics of that model will apply.

Our AI Assistant relies on Microsoft Azure OpenAI Service, specifically GPT-3.5 and GPT-4 models hosted in Azure OpenAI supporting the Chat Completion API. GPT-3.5 and GPT-4 can comprehend and generate natural language. What sets these models apart from previous models is its "conversation-in and message-out" feature, making it more intuitive and suitable for both chat and non-chat scenarios. To interact with GPT-3.5 and GPT-4, we utilize the Chat Completion API, which is a dedicated API for engaging with these models. The AI Assistant can be connected to Templafy's Azure OpenAI Service or a custom Azure OpenAI Service.

Hosted by Templafy

Azure OpenAI Service

Custom Azure OpenAI Service

GPT models hosted on Azure OpenAI Service that support the Chat Completions API can be connected – Fine-tuned GPT models are also supported since they are part of Azure OpenAI deployments (Azure AI Search/RAG -Resource Augmented Generation - and Azure AI prompt flow are not yet supported, as they are not part of the "Azure OpenAI Service" offering.

When customers use the Templafy Hosted version of Azure OpenAI, this service will be hosted in the same general location as the region selected for the customer's tenant. We guarantee that for an EU tenant, the service will be hosted in the EU region. The main regions used are Sweden Central, East US 2, Australia East and Canada East.

Templafy offers the possibility to connect to multiple AI models for the AI Assistant. As part of that, it is also possible to individually assign actions to specific connections in order to support more complex and specific use-cases per action.



6. Security and Privacy

Introduction

Security is one of the biggest concerns when discussing AI. Templafy is a security first organization, and this is reflected in our product. As organizations embrace the **transformative power** of AI, the need for a comprehensive understanding of the security implications inherent in these systems becomes imperative. We are aware of these implications, and we would like to walk you through the mechanisms that we have in place to ensure that the confidentiality and integrity of your data is maintained when using the AI Assistant.

When using AI models, we must always be aware of the risk of adversarial attacks both against our **data** and against the **model** itself.

Our AI Assistant is designed to **seamlessly integrate** with the familiar and controlled **Office environment**, prioritizing security at its core. This integration ensures that end users are encouraged to adopt this **protected iteration**, thereby safeguarding against potential risks associated with online alternatives. By fortifying security within the organization, we ensure that our AI Assistant is not only **conveniently accessible** where it is needed the most but also **aligns with the existing tech stack**, bolstering a secure and protected user experience.

Our focus is to provide users with a secure and centrally managed version of AI, which fosters an environment of **trust** and **confidence**. We believe that this strategic integration will help organizations prioritize security and protect against potential risks, thereby ensuring a seamless and productive workflow.



6. Security and Privacy

Model security

This is only applicable if the customer decides to leverage Templafy's model. If the AI Assistant is connected with a custom model, the specifics of that model will apply.

The AI Assistant is built on Microsoft Azure OpenAI Service, specifically GPT-3.5 and GPT-4 models hosted in Azure OpenAI supporting the Chat Completion API. Templafy does not have control over the underlying model, but Azure OpenAI Service has implemented several measures that raised the security level of this model to high standards. These measures include guardrails that prevent the model from generating harmful or malicious content. These guardrails include filtering out toxic or offensive content, flagging content that may be harmful and limiting the model's ability to generate certain types of content. To see more information, follow [here](#):

[Content Filter](#)

[Abuse Monitoring](#)

This latter point is also relevant for our AI Assistant. In fact, while most GenAI tools allow for free prompts, our AI Assistant limits the prompts to just the ones that were deemed acceptable by the tenant administrator. Prompt engineering can be a difficult and risky task, but with our AI Assistant you can rest assured that only pre-approved and accurate prompts are run.

Any updates and patches to the AI Assistant are deployed in accordance with our Secure Software Development Lifecycle. We secure API keys and integrations using Azure Key Vault and in accordance with best practices. More information regarding the best practices adopted across our product can be found in our latest SOC 2 report which can be requested [here: SOC 2 request report \(templafy.com\)](#).



6. Security and Privacy

Data security

Data **confidentiality** is also of the utmost importance when it comes to AI. With Azure OpenAI Service, your inputs and completion outputs, your embeddings, and your training data is **not available to other customers**.

Prompts, outputs, embeddings, and training data is **not used to improve Azure OpenAI models** or any Microsoft or third-party products or services, including OpenAI Inc.

The Azure Open AI models are **stateless**: no prompts or generations are stored in the model. Additionally, prompts and generations are not used to train, retrain, or improve the base models. Microsoft hosts the OpenAI models in Microsoft's Azure environment, and the Service does not interact with any services operated by OpenAI Inc, such as ChatGPT or the OpenAI API. There are **no third parties involved** in the processing besides Templafy and Microsoft Azure. All inputs and outputs for this service are passed directly through Microsoft OpenAI Service via secured APIs. These APIs employ various security measures such as encryption, authentication, and authorization to safeguard sensitive information from unauthorized access or tampering during transfer.

The prompts and completions data may be temporarily stored by the Microsoft Azure OpenAI Service in the same region which Templafy deploys its Azure OpenAI Service for up to 30 days for debugging purposes in the event of a failure and/or investigating patterns of abuse and misuse to determine if the service is being used in a manner that violates the applicable product terms. Human reviewers assessing potential abuse can access this repository of prompts and completions data only when that data has been flagged by the abuse monitoring system. Microsoft ensures the security of these reviewers through various means, read more [here](#).

This data is **encrypted at rest** with AES 256 and **logically separated**.



6. Security and Privacy

Data security



Templafy will improve and enhance our Service by tracking usage data. Please note that user provided content will not be tracked or logged.

The usage data is pseudonymized and includes the length in characters that users select when using the AI Assistant and the length in characters of prompts configured for actions created.



6. Security and Privacy

Privacy

The AI Assistant will process all information that is actioned by the AI Assistant. This means that if **personal data** is part of the text, this information will be processed. Processing of customer personal data is governed by the Data Processing Agreement.

Microsoft is a trusted sub-processor for delivering our services to customers. Upon the implementation of a Templafy tenant, customers select the Microsoft Azure data center regions to host the data they use for the services. Some data centers apply to the data processed for the AI Assistant, with the exception of customers who opt for data hosting in Canada (tenant is hosted in Canada Central, AI Assistant processes data in Canada East) and Europe (tenant is hosted in North/West Europe, data processed for the AI Assistant in Sweden Central). Upon an AI request, the end user's input is temporarily processed in these regions in order to generate an output, not the personal data associated to the end user (user profile information). Whether the input contains personal data depends on the end user's need and use of the AI Assistant.

As a sub-processor, Microsoft's privacy and security controls are subject to **thorough reviews** by Templafy's Information Security department, to a degree that lives up to the level of data protection depicted in the Templafy Data Processing Agreement with our customers. Both Templafy and Microsoft are certified under the **Data Privacy Framework**, the sub-processing agreement with Microsoft incorporates the **Standard Contractual Clauses**, and **Transfer Impact Assessments** for international transfers are performed.

As with all processing activities involving personal data in the Templafy services, the AI Assistant data flow complies with core data protection principles so that data is only processed for the purpose and time that have been approved and made transparent. Templafy's Privacy team has formally assessed the personal data processing by the AI Assistant and approved it as compatible with the original purpose of collecting the data, and that the risk to data subjects' rights and freedoms to privacy is not increased. That being said, **customers remain controllers** of the data used for the AI Assistant, and Templafy does not repurpose the data for its own use or other capacity than a data processor performing a contract with the customer.



6. Security and Privacy

More information

For more information regarding the overall security and privacy approach at Templafy, please refer to:

- [Product security & privacy – Templafy](#)
- [Security and Privacy FAQ at Templafy](#)
- [Data processor agreement & GDPR compliance - Templafy](#)



7. Challenges and future directions

Model and Service Hosting Locations

AI systems have the potential to revolutionize the way we live and work, but it is important to be mindful of the potential for **biases** to be absorbed from the training data. As with any learning process, AI models may unintentionally internalize biases present in the data sets used for their training. To ensure that these systems do not make discriminatory or inequitable decisions, it is crucial to approach their development with vigilance.

One of the persistent challenges in the field of artificial intelligence is the **lack of transparency** regarding the decision-making processes of AI models. These models are often considered "black boxes" that do not readily divulge the reasoning behind their decisions. It is important to advocate for increased transparency and explainability in AI algorithms to build user trust, as users seek to understand the underlying mechanisms guiding AI decisions, much like they would seek explanations from a trusted advisor.

The autonomy of AI systems highlights the need for **robust oversight**. While the autonomy of these systems is impressive, it is crucial to establish **explicit guidelines** and, where applicable, **include human oversight**. This dual-layered approach ensures that decisions, even when autonomous, adhere to established ethical standards and guidelines.

In the complex landscape of AI deployment, determining responsibility and liability for decision outcomes poses a significant ethical challenge. Establishing clear protocols for **accountability** and **liability** is essential to provide stakeholders with a transparent mechanism for recourse and redress in cases where AI systems may err or cause harm. It is important to remember that AI systems are tools to assist us and should be developed and deployed in a responsible and ethical manner.



7. Challenges and future directions

AI Governance

AI governance is crucial for ensuring **compliance, trust, and efficiency** in the development and application of AI technologies. It allows us to address the potential risks and challenges associated with AI while promoting innovation and building trust among our customers.

To achieve our objective of aligning our AI with our customers' expectations, we have taken the initiative to establish a **cross-functional AI team**. This team is responsible for developing and maintaining our AI strategy and overseeing the governance of our AI product. By having a dedicated team in place, we can ensure that our AI technologies are developed and used in a responsible and ethical manner.

Through this approach, we aim to not only meet the **regulatory requirements** but also go beyond them to build a framework that instills **trust and confidence in our AI capabilities**. We understand the importance of transparency and accountability, and our AI governance efforts are focused on creating a robust system that addresses the concerns of all stakeholders involved.

By investing in AI governance, we are not only mitigating the potential risks associated with AI but also fostering a culture of **continuous improvement and innovation**. We believe that by aligning our AI technologies with our customers' expectations, we can provide them with reliable and trustworthy solutions that meet their needs effectively.



8. Conclusion

In summary, this document aims to provide a **comprehensive understanding** of our GenAI offering. Our primary goal is to offer both technical and functional insights into this tool, including its underlying algorithms and architecture. We have also highlighted the key features available to end users and administrators, to ensure a thorough grasp of its operational aspects.

To demonstrate its **practical applications**, we have explored various use cases across organizational processes, featuring testimonials from beta testers who have experienced the tangible benefits and transformative potential of our product.

We understand the critical importance of **privacy** and **security** and have addressed these concerns transparently by discussing potential risks associated with integrating the AI Assistant into your tech stack and outlining the risk mitigation strategies that we have put in place.

Looking ahead, we are excited about the future of AI and remain committed to continuously evolving our technology to deliver even greater value to our valued customers. We hope that this document has shed light on the capabilities of our GenAI and how it can contribute to the advancement of your organization



