



AI & ANALYTICS

Generative AI Applications for the Enterprise

Business leaders across industries are captivated by the potential value of generative AI (GenAI). From automating personalized promotions to accelerating hours spent on research, the uses of GenAI to support day-to-day business operations are boundless.

Discussions of GenAI and how organizations are implementing the technology to create value are now commonplace in boardrooms and on earnings calls. However, there is currently a gap between the ambition of businesses to build their own GenAI applications and the value being realized. Gartner predicts that at least 30 percent of GenAI projects will be abandoned after the proof of concept (POC) phase,¹ resulting in significant waste of time, money and resources. To realize the potential of GenAI, businesses must consider the challenges in building enterprise-grade, production GenAI applications.

This white paper provides actionable recommendations to address these key challenges:



Select the right use case and target end users.



Understand your stakeholders and set expectations early.



Align the model with business objectives.



Uphold enterprise software development practices.



Mitigate against emerging risks.



Evolve your testing approach to incorporate new GenAI factors.

Addressing these actions will enable a strong foundation for successful application deployment and user adoption. Dive deeper into our recommendations below!

¹"Analysts Explore the Business Value of Generative AI at Gartner Data & Analytics Summit, July 29-30 in Sydney," Gartner, July 29, 2024, <https://www.gartner.com/en/newsroom/press-releases/2024-07-29-gartner-predicts-30-percent-of-generative-ai-projects-will-be-abandoned-after-proof-of-concept-by-end-of-2025#:~:text=At%20least%2030%25%20of%20generative,%2C%20according%20to%20Gartner%2C%20Inc.>



Select the Right Use Case and Target End Users

When moving GenAI applications from POCs into enterprise applications, the first critical step is selecting the right use case with the right set of end users. Chatbot applications have emerged as popular GenAI use cases, largely due to the success of OpenAI's ChatGPT product. However, this is not the only way to leverage this technology. Other valuable use cases include document and case summarization, automated communication generation, language translation and sentiment analysis. These new applications do not necessarily need to replace current processes performed by people; instead, they can enhance efficiency by helping people perform tasks more efficiently and effectively.

Starting with a small project that has well-defined goals and direct business benefits is crucial. The project must be feasible, meaning that current GenAI technologies can support the use case, the team has the necessary skills to deliver on the use case, and there is an immediate need by the business and end users for the solution. The business benefit can manifest as either a reduction in operational costs or an enhanced customer experience that drives revenue growth.

Considering the impacts of your GenAI solution is essential. How will users' decision-making and business processes be affected by the outputs of your application? How will an inaccurate or otherwise adverse model response impact customers? Organizational experience is valuable in mitigating the frequency and impact of adverse responses. For example, with an initial release of an internal knowledge base chatbot app, a team of experienced, senior associates can easily review application outputs and validate them before taking action.

This not only reduces risk and adverse customer impacts but also provides valuable feedback to the development team. A highly visible blunder can become a block to future GenAI efforts, whereas a series of smaller successful implementations can bolster the resilience of an organization's GenAI program, building the requisite skills and effectively mitigating risk.

Integrating GenAI features within existing workflows can significantly boost adoption rates. For instance, experienced researchers might favor a familiar search interface to locate relevant academic journal articles and summarize key points. In this scenario, GenAI can be seamlessly integrated into the back end, enhancing the established workflow without disrupting the user experience. This approach adds value while minimizing the need for change management. Conversely, undergraduates might prefer a chat interface to ask clarifying questions about retrieved articles. It is crucial to consider the targeted end users, their entire user journey, and the level of change management required to achieve satisfactory adoption for each use case.

Users of the application must be open to using new tools and technology. Providing the target end users with information on the purpose, background and goals of the project is critical to gain application adoption and overall project success. It is important to communicate that this technology is intended to make their jobs easier and allow them to focus on more valuable activities, such as engaging with customers.



Understand Your Stakeholders and Set Expectations Early

Early engagement and alignment with all stakeholders are critical steps in building production-ready AI applications. Key stakeholders that should be engaged include:



Business Team

End users and front-line managers who will utilize the AI application.



Development Team

Technical experts building the application.



IT Infrastructure Teams

Partner teams responsible for IT resource provisioning and application environment support.



Compliance Team

Team that ensures regulatory and policy adherence for the organization.



Business Sponsors

Leaders who fund the AI project.



IT Executives

Leaders who oversee the development team and ensure alignment with IT strategy.



Data Owners

Persons responsible for ensuring authorized data usage and compliance.



Security Team

Team that protects the organization from threats and vulnerabilities.

The alignment process begins with a clear understanding and agreement on the requirements. GenAI applications can provide a new level of flexibility that is powerful. However, at the beginning of your AI journey, this flexibility generally must be toned down to control scope, reduce timelines and limit risk.

Consider a chatbot application with an open-ended input box. This flexibility is great for ChatGPT and internet search. Your application likely does not need this level of flexibility. The open-ended, highly flexible nature of chatbots leads to development complexities and exposure to risk. It necessitates the use of robust model reasoning to accurately interpret requests and determine appropriate actions based on your predefined use cases. This requires development effort to manage complex prompt templates, routing requests to the correct logic paths, and restricting responses for unsupported use cases. A wide range of inputs must be tested to validate that the application is not responding to inputs outside its intended scope or acting on incorrect information.

The answer is not to avoid chatbots, but to design a user experience that is most effective for your use case. A poor design practice would be to implement a chat or text box UI with a lengthy list of instructions or requirements for input. This not only detracts from the application's usability but also does not prevent invalid or unintended inputs. Familiar design patterns such as forms for structured input should continue to be used in GenAI applications. End users are familiar with these patterns and can more easily learn the new application. A well-aligned approach with the business team ensures that the application remains focused, manageable and user-friendly, ultimately leading to a more successful deployment.

Business and development teams must be aligned on the tradeoffs between the benefits of a highly flexible design and the impacts on development timelines and risk.




Align the Model With Business Objectives

After user input is submitted, it is important to apply a series of validations and pre-processing steps before making a request to a GenAI model. First, identify the intent behind the input and ensure the request falls within the scope of the application. Sanitization is essential to check for and remove any personally identifiable information or harmful content, safeguarding both user privacy and the integrity of the application. Finally, input formatting is needed to ensure that the data aligns with the specific needs of the request and prompt. By implementing these measures, developers can protect the application from bad actors and inappropriate inputs, ultimately leading to a more effective and secure application.

The variety of outputs from a large language model (LLM) is to be expected and often desired. For example, travelers will likely prefer a model that produces diverse dining options for their upcoming trip rather than a system that predictably recites from the Michelin Guide.

Given the tendency for variety, it is important to align with the business team on the range of acceptable outputs and where the variability needs to be controlled. Receiving a structured output, like JSON, from the model simplifies the process for applying application logic and subsequent processing of the content. Based on the model response, application flows may need to be redirected, for example, to get additional information or clarity from the end user. By implementing such measures, the application can better handle a wide range of outputs effectively, ensuring that end users receive consistent and reliable responses. The model should be focused on its reasoning capabilities and generating natural language text or images, leaving input and output data formatting to the application logic, where there is more control.



Uphold Enterprise Software Development Practices

Many businesses approach the development of GenAI applications differently from traditional software projects, often treating them as exploratory R&D efforts. This approach can result in insufficient resource support, such as dedicated staff, sustained funding, appropriate IT infrastructure, or partner team collaboration. Consequently, these projects often cut corners when it comes to planning, risk mitigation and security.

Consider an organization that is creating a GenAI application proof of concept (POC). Because the application is a “POC,” the organization assigns the application to an “R&D” development environment that lacks the features and service level agreements (SLAs) of a traditional environment. This approach may provide infrastructure resources at a lower cost, but it can also lead to delays and higher development costs on the application side. It is common for less-supported environments to experience higher levels of downtime, which burns development and testing cycles. Further, “R&D” environments may lack security and monitoring tools, which increase the risk of adverse events such as noncompliant data use. Lastly, infrastructure management processes and personnel may be unavailable, which increase the likelihood of “rogue IT.”


Infrastructure configurations that do not follow established organizational patterns are often less effective, less secure and less maintainable. If the pilot manages to surmount the many challenges stemming from the “R&D” environment, it may hit additional challenges when going live and meeting production level requirements, leading to re-work and accrued technical debt.

It is shortsighted to bypass best practices such as version control, unit tests and code reviews in the pursuit of rapid R&D. Instead, a more effective approach is to consider GenAI as another tool or service available to the development team. By applying the organization’s existing software development processes and standards to GenAI projects, the team can ensure robust and reliable outcomes. This includes maintaining best practices for version control, code reviews, and comprehensive planning for integration, security and compliance.

Staffing GenAI projects similarly to traditional software projects allows the team to benefit from a broad array of necessary skills. While there is an additional need for specialized GenAI skills, the importance of other roles such as front-end engineering, back-end engineering, DevOps, data engineering and product management remains crucial. These roles ensure that the project is well-rounded and capable of addressing all aspects of development, from user interface design to infrastructure management.

GenAI applications, in particular, introduce a range of technical considerations due to their reliance on various tools throughout the stack, including code frameworks, model vendors, compute platforms and data storage services. Like all other technologies used within an organization, these AI services should undergo thorough evaluation and approval processes to ensure robust security and compliance.

This includes the assessment of GenAI models to ensure they meet the organization's security and compliance standards. Organizations benefit from consistent use of approved tooling, as repeated experience with familiar tools can streamline development processes and enhance security measures. However, the rapidly evolving landscape of GenAI necessitates a balance between leveraging established tools and exploring new technologies that offer improved capabilities. Teams must continuously evaluate the “tried and true” against the “latest and greatest” to maintain a competitive edge while ensuring security and compliance.



Teams must continuously evaluate the “tried and true” against the “latest and greatest” to maintain a competitive edge while ensuring security and compliance.



Mitigate Against Emerging Risks

GenAI applications are subject to privacy, compliance and security risks that are faced by traditional applications as well as new threats such as biased training data and prompt manipulation. These emerging threats are discussed in a recent A&M whitepaper.² These risks underscore the need for rigorous security reviews to prevent and mitigate potential attacks from malicious actors. Applications that handle sensitive data or facilitate critical business processes are particularly valuable targets, making their security paramount.

Security teams must analyze the types of data provided to GenAI models and the methods by which these models process that data. It is advisable to avoid including sensitive data, such as personally identifiable information (PII) and protected health information (PHI), unless necessary for the use case. The inclusion of such data not only imposes additional compliance obligations but also increases the risk of introducing unwanted biases into the model's responses.

Additional security guidelines to consider are:



Review of the geographic residence and infrastructure ownership of the selected GenAI model.



Log and regularly audit model inputs and responses.



Routinely test the application against attack and data leakage scenarios.



Automate data cleanup to ensure data retention policy compliance.



When leveraging RAG, implement compliant role-based access controls for data elements and sources to control user access to data.



Inspect and validate user inputs for injection attacks and model misuse.



Protect application endpoints by limiting repeated requests (rate limiting), especially for LLM calls that can generate significant vendor cost.

Planning and building support up front for these security and compliance features will ultimately save time and reduce risk in your project. These practices promote reliability and trustworthiness, ultimately contributing to the successful deployment and adoption of AI technologies within the organization.

²"Training Data and Prompt Manipulation: How to Keep Your Organization Safe Against Large Language Model Cyberthreats," Alvarez & Marsal, June 12, 2024, <https://www.alvarezandmarsal.com/insights/training-data-and-prompt-manipulation-how-keep-your-organization-safe-against-llm>.



Evolve Your Testing Approach to Incorporate New GenAI Factors

In addition to utilizing standard software development processes, GenAI application development teams will need to apply more effort to application testing and AI ethics safeguards. The variability introduced by GenAI models must be accounted for in the testing phase. Traditional software systems are deterministic, meaning that given inputs A and B, the output will always be X. In contrast, GenAI applications often produce non-deterministic outputs, where the same inputs A and B can yield varying proportions of outputs X, Y and Z. This variability underscores the importance of evolving testing processes to include extensive regression testing and repeated testing of common user journeys. By maintaining detailed test documentation, product owners can gain valuable insights into the distribution of possible outputs and ensure the reliability of the application.

It is crucial for organizations to implement ethical guidelines and frameworks to ensure responsible AI development and deployment. AI ethics considerations focus on fairness, transparency and accountability. Without proper safeguards, organizations risk violating privacy rights, perpetuating biases, or creating systems that are difficult to control. Consider, for example, an organization that releases a GenAI photo editing tool that underperforms on certain skin tones. This outcome is ethically unacceptable. Further, the organization may face business and reputational harm, as users feel alienated by the application. These types of situations can be minimized by robust testing and alignment with the business team on the variety and distribution of application responses.

For the underperforming photo editing tool, these types of issues are often caused by the GenAI model using training data sets which are not sufficiently diverse. This can be rooted in the base model or a result of fine tuning. Proactively collecting diverse training data can help mitigate these shortcomings. Further, teams should not simply consider an “overall accuracy” metric. Instead, they should create and track group-level evaluation metrics to ensure equitable performance. After establishing these metrics, the team can iteratively discover and tackle performance discrepancies to reduce ethical risk before releasing the application. For more on AI ethical risks, consult a recent A&M white paper, “AI Ethics Part One: Navigating Pressures for Responsible AI,”³ which delves into these challenges in greater detail.

Best practice is to establish your organization’s overarching AI strategy. A mature AI strategy outlines the do’s and don’ts for building AI applications, providing a roadmap for development teams to follow. Adhering to these guidelines ensures that GenAI projects are consistent with the organization’s goals and values.

³“AI Ethics Part One: Navigating Pressures for Responsible AI,” Alvarez & Marsal, August 20, 2024, <https://www.alvarezandmarsal.com/insights/ai-ethics-part-one-navigating-pressures-responsible-ai>.

Conclusion

Building production-ready GenAI applications requires a comprehensive approach that includes alignment with stakeholders, proper IT support, standard software development processes, and new AI-specific considerations. Organizations must spend additional effort on aligning expectations with business owners, managing model inputs and outputs, and mitigating security and compliance risks. By integrating these elements into the organization's AI strategy, organizations can maximize the potential of GenAI while maintaining high standards of reliability, security and ethical responsibility. With successful deployments and end user adoption, the organization can finally experience the much-anticipated value that everyone has been looking forward to.

AUTHORS



David Dina
Senior Director
Atlanta, GA
ddina@alvarezandmarsal.com



Gray Cannon
Manager
Atlanta, GA
gcannon@alvarezandmarsal.com

ABOUT ALVAREZ & MARSAL

Follow A&M on:

Founded in 1983, Alvarez & Marsal is a leading global professional services firm. Renowned for its leadership, action and results, Alvarez & Marsal provides advisory, business performance improvement and turnaround management services, delivering practical solutions to address clients' unique challenges. With a world-wide network of experienced operators, world-class consultants, former regulators and industry authorities, Alvarez & Marsal helps corporates, boards, private equity firms, law firms and government agencies drive transformation, mitigate risk and unlock value at every stage of growth.

To learn more, visit: [AlvarezandMarsal.com](https://www.alvarezandmarsal.com)

© Copyright 2025 Alvarez & Marsal Holdings, LLC.
All Rights Reserved.
460766-46915/March 2025
9561_STG02