SIAPARTNERS

Understanding Generative AI:

Insights for Safe Tool Selection

oct 2023

Summary

- Reviews critical factors like data privacy, documentation, legal compliance, model validation, and IP concerns when selecting a Generative AI tool.
- Discusses building versus buying considerations, including organizational needs, capabilities, costs, risks, and strategic alignment.
- Emphasizes responsible and ethical AI practices aligned with regional regulations.
- Stresses thoughtful decision-making based on diligent vendor evaluation and understanding of opportunities and risks.
- Aims to provide a comprehensive guide to safely acquiring and implementing Generative AI.

Are you ready to navigate the complex world of Generative AI (GenAI), but worried about the risks ?

In today's fast-paced technological landscape, GenAI offers immense transformational opportunities across industries, but it also brings complexity. Selecting whether to build or buy a GenAI tool is a process filled with challenges and requires a detailed understanding that intersects between Compliance, Risk Management, Legal, and Technology/Analytics.

This article introduces a roadmap, highlighting key factors to consider before choosing a GenAI solution. From data privacy nuances to intellectual property considerations, it guides you through the maze of making an informed decision that aligns with your unique needs and goals.

While this guide is primarily focused on enterprise software, its insights and recommendations are mostly applicable to any B2C usage, including individuals buying a tool for their own use.

The critical aspect of building versus buying is explored, emphasizing the balance between your organization's individual requirements, abilities, and long-term vision. It's about thoughtful choices, not just embracing what's new and exciting.

Whether you are a business leader, tech enthusiast, or curious explorer of GenAI, this article arms you with the essential understanding needed to make the right choice in this intricate and rapidly evolving field.

Before we deep dive into each topic, here is an overview of key factors to review before getting started:

• ***Data Privacy***	• *Evaluate the vendor's privacy practices, covering aspects such as data minimization, anonymization, consent, security, subprocessor, retention, disposal, updates, compliance with laws, and a commitment to "Privacy by Design".*
• ***Model Validation x Audit***	• *Ensuring that the GenAI tool operates within legal, fair, and accurate bounds is crucial. Look for vendors who offer independent validation/audit reports, have a comprehensive scope covering data privacy, soundness, and bias, and align with regional regulations like the EU's AI Act, GDPR, or the U.S.'s SR 11-07 guidelines. The tool should also be auditable.*
• ***ESG x Ethics and Compliance***	• *A well-rounded GenAI tool should align with ESG goals, from energy-efficient hardware and model optimization to ethical data practices and inclusivity. Evaluate the vendor's adherence to global regulations like the EU's Ecodesign Directive or the U.S.'s Federal Data Center Optimization Initiative to ensure the tool meets comprehensive environmental and societal standards.*
• ***Intellectual Property***	• *Ensuring compliance with IP laws is vital when selecting a GenAI tool. Focus on data licensing and permissions, be cautious of data scraping practices, and consult legal advice for jurisdictional compliance. Vendor diligence, including checking for proper IP adherence and requesting transparent documentation, is key to avoiding legal risks.*
• ***Documentation***	• *Documentation must strike a balance between transparency and proprietary protection, guided by the EU AI Act's principles of safety, transparency, accountability, and rights protection. Specific areas to scrutinize include data description, model architecture, evaluation metrics, and known limitations to ensure alignment with legal compliance and specific needs.*

Data Privacy

Read the Privacy Policy and your Data Protection Addendum ("DPA")
Don't sign up your company or customer's data to a GenAI tool without thoroughly evaluating privacy practices. Here are some critical aspects to consider:

Data Minimization ("Only collect what is needed"): Verify that the vendor collects and uses only the data essential for the AI model's intended purpose. They should explicitly avoid gathering unnecessary or overly sensitive information, aligning with privacy best practices.

Anonymization and De-identification: Look for vendors who actively remove or modify personally identifiable information (PII) within their datasets, a crucial step in preventing individuals from being directly identified. Also, evaluate the algorithms used and whether they can be decrypted/reversed. Another approach is to use aggregated datasets.

Consent and Transparency: Ensure that the vendor follows transparent data processing practices and adheres to informed consent principles if they collect data directly from individuals.

Data Security: Examine the security measures implemented by the vendor to protect against unauthorized access, breaches, or leaks. Robust measures like encryption, access controls, and regular audits are vital. You should request the vendor provide an architecture diagram, clearly showing how the data flows between various systems/vendors, which data is being sent in/out, for which purpose, and the retention period.

Subprocessors / Fourth Parties: Ensure clarity on fourth parties (vendor's third parties) with access to your data, including their specific usage and countries of operation.

Data Retention: The vendor's policy should clearly define the length of time data is stored, the conditions under which it is retained or deleted, and the security measures in place to protect it. Lastly, transparent communication with data subjects about retention practices is vital.

Updates: Oftentimes, the vendor will be able to update their privacy policy with a notification process to the customer, make sure you review those updates whenever they are published. Pro Tip: ask your favorite GenAI tool to summarize the new obligations.

Data Deletion Requests: You should review how the vendor will meet data deletion requests, especially if they embed personal data within their training datasets.

Legal and Compliance: Stay informed about the vendor's adherence to relevant data protection laws like GDPR, CCPA, HIPAA, and others, depending on your jurisdiction. Compliance with these regulations is non-negotiable.

Vendor Terms and Agreements: Review the vendor's terms of service, privacy policy, and Data Processing Agreement (DPA) if signing a separate one. Make sure you review any (hidden) referenced website as part of your agreement.

Privacy by Design: Lastly, this is more subjective but you should look for a commitment to integrating privacy considerations throughout the AI model's life cycle, and in their processes. This "Privacy by Design" approach reflects a deep, foundational commitment to privacy protection. For example, the vendor should be ready to go through a comprehensive review before you sign up with the tool.

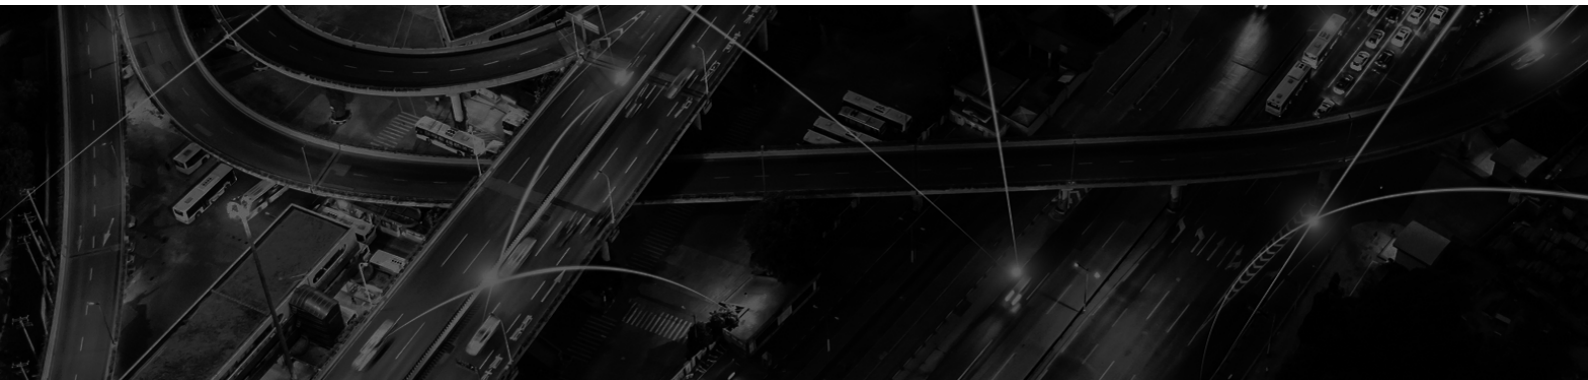

Model Validation x Audit

Model validation is a critical factor in selecting a Generative AI tool, as it ensures that the AI system operates fairly, accurately, and within legal boundaries. Here's what you should consider when evaluating a GenAI tool for model validation:

Independent Report: An independent model validation report should be provided or performed before going live with a Generative AI solution, especially if it deals with sensitive data.

Scope: The scope of the model validation/audit should include not only data privacy and security but also an evaluation/testing of the model, its soundness and design choices, as well as any relevant documentation.

Validating Models for Bias: Look for vendors that prioritize validation to maintain accuracy and ensure models don't discriminate. Different laws and regulations may include requirements to assess AI models for fairness, accuracy, non-discrimination, bias, and consumer risk. A clear validation methodology is key.

Monitoring: Consider the vendor's approach to monitoring their models in production. Whether in real-time or at regular intervals, this monitoring helps in responding to unexpected or dangerous incidents. Ensure their strategy aligns with your organization's expectations and capabilities.

Regional Considerations: Based on the type of institution you are in (e.g., financial institution), your region, as well as the type of data you're handling, you should review requirements set out by your local regulators. Here is a selection of relevant regulatory considerations.

European Union:

- *AI Act (Proposed): Check for compliance with the EU's proposed AI Act, focusing on high-risk AI systems, including mandatory validation processes.*
- *GDPR: Ensure adherence to GDPR's transparency, accountability, and data protection principles, particularly Article 22 concerning automated decision-making.*
- *European AI Alliance's Ethics Guidelines: Look for alignment with these guidelines that emphasize validation, fairness, transparency, and accountability throughout the AI lifecycle.*

Asia:

- *Model AI Governance Framework: If relevant to your region, evaluate the vendor's alignment with Singapore's framework, encouraging continuous model validation, transparency, and accountability.*
- *Model validation is not just a technical consideration but a comprehensive evaluation involving legal, ethical, and regional factors. A vendor's commitment to robust model validation practices reflects a broader understanding of the complexities of AI in a global, interconnected society.*

United States:

- *Under SR 11-07, the Federal Reserve has published guidelines on how to define, document, monitor and review/validate models on an ongoing basis.*

Environmental, Social, and Governance (ESG) & **Ethics and Compliance**

Considering Environmental, Social, and Governance (ESG) factors while buying an AI tool is a responsible and forward-thinking approach.

As AI continues to proliferate, its ESG implications are becoming increasingly significant. Responsible AI development requires that these considerations be integrated at all stages, from conception to deployment.

1. Energy Efficiency: AI tools often require significant coputational power. Consider the following to be more energy-efficient:

Hardware Selection: Utilizing energy-efficient hardware like low-power processors, GPUs, TPUs, and optimizing energy-efficient data centers can significantly minimize energy consumption.

Model Optimization: Creating more efficient models by minimizing computational complexity and memory requirements can lower energy consumption without sacrificing performance.

Data Efficiency: Reducing redundant data and focusing on data quality can minimize the energy needed for data processing and storage.

2. Ethical Practices: Choose Ethical AI, which aims to create a positive societal impact, promotes advancements that benefit communities and individuals, and strives for a responsible integration of AI into society. Take the following into consideration:

Ethical Data Usage: Sustainable and ethical data sourcing and avoiding harmful practices like excessive data collection are vital to responsible AI development.
Ethical alignment with your organization's values is key as well. Ask about biases, fairness concerns, or any ethical considerations as they relate to the AI model's development and deployment, in line with the EU AI Act's emphasis on ethical AI practices.

3. Inclusivity and Fairness: Inclusivity in AI development involves ensuring that AI systems are accessible and beneficial to a diverse range of users, regardless of their backgrounds, abilities, or demographics. Fairness refers to the equitable treatment of individuals and groups, ensuring that the outcomes of AI systems are not biased against or favouring any group. The following should be kept in mind while developing/choosing an AI tool:

Diverse and Inclusive Models: Ensuring training data is diverse and representative helps avoid bias and discrimination and provides an equitable outcome for all.

Monitoring and Testing: Regular evaluations and real-world monitoring of AI models are critical and ensure that they remain unbiased.

Human-Centric Design: Creating AI products that are accessible to individuals of all ages, genders, and abilities promotes inclusivity. This includes considering the needs of marginalized or underrepresented groups. AI developers should be aware of historical biases present in data and strive to avoid perpetuating these biases in AI systems.

4. Global Regulations: Several global regulations encourage considerations of ESG factors:

European Union: The **Ecodesign Directive** and **Energy Efficiency** Directive guide energy-related products, encouraging energy-efficient technologies. The proposed **AI Act** also includes provisions for high-risk AI systems to undergo conformity assessments, including considerations for environmental impacts.

North America: The U.S. emphasizes green procurement through initiatives like the **Federal Data Center Optimization Initiative**, pushing for more energy-efficient AI technologies.

Asia: Countries like China and Japan are establishing standards impacting AI hardware and promotion of conservation through the **Energy Efficiency Standards** and **Energy Conservation Act** respectively.

The ESG considerations in AI are broad and multifaceted. A comprehensive approach encompassing energy efficiency, ethical practices, inclusivity, and adherence to global regulations ensures that AI development aligns with broader societal and environmental goals. The push for responsible AI not only enhances its societal acceptance but also fosters innovation in a way that respects our planet and its diverse inhabitants.

Intellectual Property (IP) Concerns

Intellectual Property (IP) rights play a critical role in acquiring and deploying AI tools. Violations of patents, copyrights, trademarks, or trade secrets can result in significant legal liabilities and have led to recent lawsuits, particularly concerning the unauthorized use of scraped data. Here are the key considerations when acquiring an AI tool, and recognizing the variations in IP rules across jurisdictions:

1. Data Licensing and Permissions:
You should always check with the Vendor about the origin of the data used for training the AI tool. Four types of data sources will place you on the safe side:

- **Authorized Data:** Always prioritize AI tools developed using datasets released under open licenses or publicly available data initiatives.
- **Data from their own collection:** It is also safe to use AI tools trained using data collected or simulated by the developer. Be careful, you still need to remember all the privacy requirements above.
- **License Agreements:** If the developer has permission to use copyrighted or proprietary databases, your project can get the green light.
- **Data obtained through synthesis techniques:** Your vendor can safely use synthetic data that has been generated artificially and that mirrors the original data's statistical characteristics.

2. Data Scraping:
AI developers often use data-scraping methods to obtain training data. This method implies that they collect information from various sources on the internet without obtaining proper authorization from the website owners. This technique comes with multiple legal and ethical risks that you need to be aware of, especially given recent lawsuits in the United States.

- **Violation of website terms of use:** Some websites explicitly prohibit scraping in their terms of use, while others impose specific guidelines on using their content.
- **Violation of privacy rules:** There is a high chance that scraped data contains personal or other sensitive information. Its use might be against the law.
- **Violation of intellectual property rights:** Data scraping can lead to IP law violations if it involves unauthorized copying or use of protected content, such as copyright or trademark material. Multiple Generative AI developers are being sued for violations of copyright law in a series of lawsuits in the United States.

3. Jurisdictional Compliance:
Understand Regional IP Rules: IP rules vary by jurisdiction, so it is crucial to comprehend the laws of the territory where the AI tool will be deployed.

Legal Consultation: The good news is that copyright materials could be used in certain situations to train AI models. Seek legal advice for possible exemptions and specific compliance with copyright rules in each jurisdiction.

4. Upcoming Regulations:
AI Act (European Union): The proposed EU legislation may require Generative AI developers to publicly document the use of training data protected under copyright law.

China: The recently adopted Interim Measures for the Management of Generative Artificial Intelligence Services clearly states that Generative AI service providers must refrain from infringing upon the intellectual property rights enjoyed by others.

5. Vendor Diligence:
As you review your vendor, keep in mind the two items below:

- **Vendor IP Compliance:** Ensure that the AI tool provider adheres to all IP laws and has obtained the necessary licenses for data and technology used.
- **Transparency and Documentation:** Request detailed information about the AI tool's development process, including data sourcing, to avoid potential IP issues later.

In the age of data-driven technology, IP considerations are paramount when purchasing an AI tool. Organizations must adopt a diligent and well-informed approach, recognizing the complexities of IP laws and the consequences of non-compliance. Attention to data licensing, jurisdictional rules, upcoming regulations, and rigorous vendor examination can guide a lawful and ethically sound acquisition.

Documentation

In the complex world of Generative AI, transparency and understanding are key. Vendors, rightfully, may not disclose every aspect of their proprietary algorithms. However, a certain amount of documentation should be made available – aligning with the recent guidelines published by the European Union in the draft AI Act, as well as best practices shared by regulators around model risk management.
The draft EU AI Act is a comprehensive legal framework proposed by the European Union to govern AI systems, and seeks to:

Ensure Safety: By setting requirements for high-risk AI systems, the Act promotes safety and consistent performance across member states.

Promote Transparency: The Act emphasizes transparent data handling and reporting, requiring vendors to disclose general information on datasets, algorithms, and usage even when specific details are proprietary.

Facilitate Accountability: It sets up structures for continuous oversight, monitoring, and regulatory enforcement to ensure compliance with safety and ethical standards.

Protect Fundamental Rights: The Act safeguards individuals' rights, ensuring that AI systems do not lead to discrimination or other unfair practices.
Here's what to look for within the bounds of what is typically shared:

Data Description: Vendors should offer an overview of datasets used, sources, preprocessing, and compliance with transparency requirements such as the EU AI Act. Review how usage data influences the re-training or tuning of algorithms.

Model Architecture and Algorithms: While exact algorithms may be proprietary, a high-level description of the model's architecture, type, and key design choices should be available, or available verbally after discussion with the vendor. Look for insight into the rationale behind these choices to gauge the tool's effectiveness.

Evaluation Metrics: The vendor should share the metrics used to assess the model's performance, such as accuracy and precision. This allows for an objective evaluation, even if specific training details are proprietary.

Training Data Sets: Understanding the sources, collection, and preparation of training data is vital. Ensure the vendor's process aligns with applicable privacy laws and your organization's values and requirements, as you don't necessarily want the vendor collecting your prompts that may include confidential or customer information to train their models.

Known Issues and Limitations: Look for an honest assessment of the model's limitations and challenges. While vendors may not reveal every detail, acknowledging known issues shows integrity and will help you anticipate potential challenges.

Choosing the right Generative AI tool requires a delicate balance between understanding what makes it tick and respecting the proprietary nature of the technology. The EU AI Act adds another layer of complexity, stipulating certain transparency and ethical standards that must be met. By focusing on these aspects, even without access to every single step of the algorithm, you can gain valuable insights into a tool's capabilities, ethics, and legal compliance, leading to an informed decision that aligns with your specific needs and regulations.

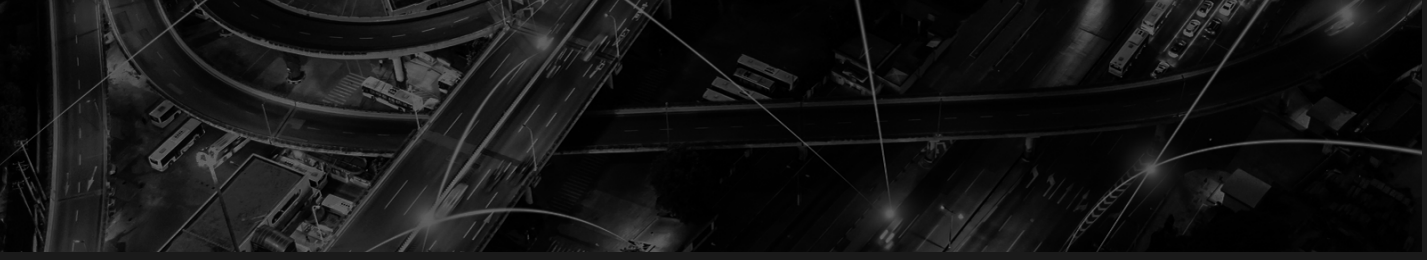

Building vs. Buying a Generative AI Tool

After reading these considerations, you must wonder, can you build all of this yourself instead of relying on a vendor? This decision is far from trivial and requires a nuanced understanding of various factors. Here's a guide to help you navigate this complex decision-making process:

1. Understanding Your Needs and Objectives

- **Specific Requirements:** Define and document the unique needs and objectives of your organization. Do you require specific features that are not available in existing tools?
- **Strategic Alignment:** How does the GenAI tool align with your long-term business strategy? Building a solution might provide a competitive advantage if it's highly specialized.
- **Document Success Metrics:** What does success look like? Make sure you determine success metrics from the beginning, whether it's time/money saved, improved compliance or reliability, or time spent on less manual tasks.

2. Evaluating Internal Capabilities

- **Technical Expertise:** Do you have the in-house expertise to build, maintain, and upgrade a GenAI solution? Assessing your internal capabilities is crucial.
- **Resource Allocation:** A solution may require significant time and financial investment, as well as expertise. Evaluate your budget and willingness to commit resources, potentially medium-long term.

3. Analyzing Market Options

- **Vendor Evaluation:** Assess existing GenAI tools in the market. Can they fulfil your requirements without customization?
- **Cost Comparison:** Compare the total cost of ownership (TCO) for both building and buying. Include ongoing maintenance, support, and potential upgrades.

4. Compliance and Risk Management

- **Legal Compliance:** Are there specific legal or regulatory considerations that may affect your decision to build or buy? What are the internal processes that you need to get approval from before you start this journey (this could include supplier security assessment, penetration testing, architecture review, and code review, based on your company policies)
- **Risk Assessment:** Analyze potential security, privacy, and intellectual property rights risks.

5. Scalability and Future Growth

- **Adaptability:** How easily can the solution adapt to future changes in technology or business needs?
- **Long-term Viability:** Consider the long-term sustainability of the chosen approach. Will building provide more control and flexibility, or will buying offer better support and quicker access to new features?

6. Ethical and Cultural Alignment

- **Organizational Culture:** Reflect on how building or buying aligns with your organizational culture and values.
- **Social Responsibility**: Consider ethical implications, such as environmental impact, responsible AI practices, and community engagement.

Making the Informed Choice

The decision to build or buy a Generative AI tool is multifaceted, involving careful evaluation of your organization's unique needs, capabilities, and strategic goals. Neither approach is inherently superior; the best choice depends on various factors specific to your situation.

Engage with stakeholders, conduct thorough research, and utilize expert consultations if needed. Consider developing a decision matrix to weigh the pros and cons objectively. Remember, the goal is not merely to embrace the latest technology but to make thoughtful and informed decisions that align with your organization's core values and long-term vision.

SIA GPT

Conclusion

In the rapidly evolving landscape of Generative AI (GenAI), making an informed tool selection isn't merely a technical decision—it's a multifaceted commitment to ethical, legal, and societal responsibilities. «Understanding GenAI: Insights for Safe Tool Selection» aimed to provide a comprehensive guide across critical areas, including Data Privacy, Documentation, Model Validation x Audit, ESG considerations, and Intellectual Property (IP).

Navigating this complex maze requires a balanced approach, combining technical proficiency with ethical responsibility and legal compliance. This is where Sia Partners can add significant value. We understand these nuances as a global management consulting firm with deep expertise at the intersection of Compliance, Risk Management, Legal, and Technology/Analytics. We've even developed our own GenAI solution, SiaGPT, which specializes in multi-document data extraction using multi-LLM GenAI, showcasing our hands-on experience in this space.

The journey toward selecting the right GenAI tool is fraught with challenges but also ripe with opportunities for organizations to innovate responsibly. Whether you're looking to build or buy, Sia Partners can guide you through this journey, ensuring that the GenAI tools you select are technically robust, ethically sound, legally compliant, and aligned with your overarching business and societal goals.

By making informed decisions across these vital areas, organizations can ensure they are responsible stewards of powerful technology, ultimately enhancing consumer trust and confidence in GenAI solutions.

SiaGPT

SiaGPT stands out as an AI-centric tool tailored for businesses that prioritize precision and efficiency in data interrogation from expansive databases. Its chatbot interface is not only interactive but also deeply rooted in transparency—users can pose questions, receive instant answers, and have the power to view and highlight the exact source within the document. SiaGPT's capabilities stretch further to encompass batch data extraction and an insightful scorecard function, enabling users to dissect multiple data points across numerous documents. As a nod to its commitment to security, SiaGPT ensures no training on user prompts, and its multi-LLM feature allows businesses to deploy LLM on their servers, bypassing the constraints of APIs. The tool's multilingual capabilities cater to a global user base. Above all, its intuitive design means that businesses can hit the ground running, making data management and utilization not only effective but also effortless.

Your contact

Cyril **SAYADA**

Managing Director
Legal, Risk & Compliance
San Francisco, California
Cyril.sayada@sia-partners.com
+1 (929) 363-9791

SIAPARTNERS

About Sia Partners

Sia Partners is a next-generation management consulting firm and pioneer of Consulting 4.0. We offer a unique blend of AI and design capabilities, augmenting traditional consulting to deliver superior value to our clients. With expertise in more than 30 sectors and services, we optimize client projects worldwide. Through our Consulting for Good approach, we strive for next-level impact by developing innovative CSR solutions for our clients, making sustainability a lever for profitable transformation.

www.sia-partners.com