


NAVIGATE YOUR GENERATIVE AI JOURNEY

STEP BY STEP GUIDE



By Claudia Mejía
May, 2024
www.ikigai-edge.com





1

IDENTIFY RELEVANT BUSINESS USE CASES




- **Focus on Future Potential:** Move beyond basic automation and seek opportunities where Generative AI can create entirely new possibilities for your company.
- **Collaborate with Business Leaders:** Work closely with various teams to understand their pain points, strategic goals, and data insights.
- **Think Big (But Start Small):** Brainstorm innovative use cases while keeping the initial pilot project manageable and focused.

Refer to Appendix for Business Cases Examples



- **Replicate Existing Processes:** Avoid simply replicating current tasks with AI. Look for ways to fundamentally improve or even reinvent workflows.
- **Focus on Technical Feasibility Alone:** While technical aspects are important, prioritize the strategic impact and business value.
- **Go Big Bang:** Start with a well-defined pilot project to test the waters and learn before scaling up.



2

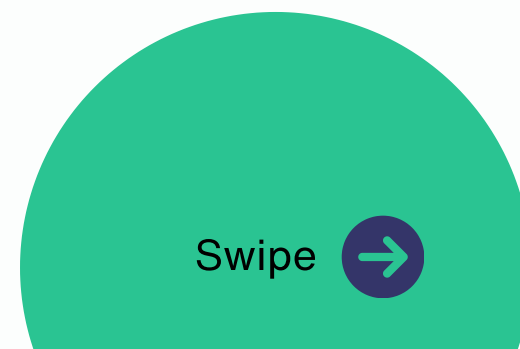
PRIORITIZE THE USE CASES FOR YOUR PILOT



- **Match Use Cases to Strategic Goals:** Select cases that directly address key business objectives identified in the previous step.
- **Consider Feasibility and Impact:** Choose a use case that is technically achievable within your pilot timeframe but also has the potential to deliver significant results.
- **Favor Learning Over Perfection:** Prioritize a use case that offers the most learning opportunities, even if it's not the most "polished" option.



-
- **Prioritize Popularity:** Don't simply choose the most popular use case for generative AI. Focus on what best suits your company's specific needs.
 - **Underestimate Data Requirements:** Ensure you have access to the necessary data sets to successfully train and test your chosen use case.
 - **Overlook Change Management:** Consider the potential impact on employees and factor in change management strategies for a smooth pilot implementation.





3


ASSEMBLE A SMALL BUT DIVERSE TEAM




- **Assemble a Cross-Functional Team:** Include business experts familiar with the impacted areas, data scientists for model customization, AI/ML engineers for integration, and IT support for maintenance.
- **Emphasize Diversity of Thought:** Seek individuals with varied backgrounds and perspectives to spark creative solutions and ensure well-rounded decision-making.
- **Empower a Project Manager:** Include an experienced project manager to keep the pilot on track, manage resources, and facilitate communication between team members.



-
- **Focus on Quality Over Quantity:** Prioritize a lean, efficient team with the right skillsets, rather than a large group with overlapping responsibilities.
 - **Neglect Governance:** Establish a dedicated generative AI governance team early on to define responsible use policies and ensure ethical implementation within the pilot.
 - **Overlook Team Dynamics:** Cultivate a collaborative environment where team members can openly share ideas and work towards a common goal.





4

DESIGN AND PLAN THE PILOT



- **Define Success with KPIs:** Establish clear and measurable Key Performance Indicators (KPIs) tailored to your chosen use case. This allows you to track progress and assess the effectiveness of the AI pilot.
- **Embrace the "What-IFs":** Proactively identify potential risks, including model hallucinations (generating inaccurate results), biased outputs, data privacy concerns, and security vulnerabilities.
- **Develop Mitigation Plans:** For each identified risk, create a mitigation plan outlining strategies to minimize its impact.



-
- **Fly Blind Without KPIs:** Don't launch the pilot without specific metrics to gauge its success. Without measurable goals, it's difficult to assess AI true value.
 - **Ignore Ethical Considerations:** Don't underestimate the potential for bias, privacy concerns, or security risks. Integrate ethical considerations into the design phase to ensure responsible AI development.
 - **Leave Risk Management to Chance:** Don't wait for problems to arise. Plan for potential pitfalls and develop mitigation strategies in advance for a smoother pilot experience.



5

DECIDE ON THE DEPLOYMENT APPROACH



- **Consider Use Case Needs:** Align the deployment approach with the specific requirements of your chosen use case. (Refer to [appendix for each approach](#))
- **Think Scalability:** Factor in potential future growth when making your decision. If you envision broader adoption within the company, choose an approach that can scale effectively. (Consider cloud-based solutions for scalability)
- **Evaluate Your Resources:** Be realistic about your in-house expertise and available resources. Opt for a pre-built solution (Consume or Embed approach) if your team lacks experience in building and maintaining complex AI models.



-
- **Overlook Security:** Ensure the chosen deployment approach adheres to your organization's data security standards. Cloud-based solutions might require additional security considerations.
 - **Underestimate Integration Complexity:** Don't underestimate the effort required to integrate a custom model (Build approach) into your existing workflows. Evaluate the technical expertise needed for a smooth integration process.
 - **Neglect Vendor Lock-in:** If opting for a pre-built solution (Consume or Embed approach), be mindful of potential vendor lock-in, where you become reliant on a specific vendor's platform.



6

DECIDE ON THE DESIGN PATTERNS



- **Align Pattern with Deployment:** Choose a design pattern that complements your chosen deployment approach ("Consume," "Embed," or "Build") - **Refer to Appendix for Patterns**
- **Prioritize Use Case Needs:** Select a pattern that best addresses the specific requirements and complexities of your chosen generative AI use case.
- **Consider Expertise:** Be realistic about your team's capabilities. Opt for simpler patterns like "LLM As-Is" or "Embed LLM" if building complex models (e.g., "Fine-Tuning LLM") is not feasible.




-
- **Overlook Flexibility:** While a specific pattern is chosen initially, consider potential future modifications and choose a pattern that allows for **some adaptation as your use case evolves.**
 - **Focus Solely on Complexity:** Don't get caught up in choosing the most advanced pattern. Sometimes, a simpler pattern like "LLM with Document Retrieval" can deliver valuable results for your pilot.
 - **Neglect Scalability:** If you envision broader adoption within your company in the future, factor in scalability when choosing a design pattern.



7

ASSESS DATA INFRASTRUCTURE

- **Identify Data Requirements:** Clearly define the specific data types and formats needed to train and run your chosen generative AI model effectively. Refer back to the use case and design pattern for specifics.
 - **Inventory Existing Data Sources:** Take stock of your existing data repositories like databases or customer records, Identify data that aligns with the pilot's needs.
 - **Evaluate Data Quality:** Ensure the identified data sets are clean, accurate, and free of biases that could negatively impact the AI model's performance. This might involve data cleansing or enrichment procedures.
-
- **Underestimate Data Needs:** Don't underestimate the amount and quality of data required for successful AI training. Inadequate data can lead to poor model performance and inaccurate results.
 - **Overlook Data Governance:** Be mindful of data privacy regulations and internal data governance policies. Ensure you have the necessary permissions and adhere to ethical data practices when using data sets.
 - **Neglect Data Preparation:** Don't simply dump raw data into the AI model. Invest in cleaning, pre-processing, and potentially enriching your data to optimize the model's learning process.





8


CHOOSE THE RIGHT HARDWARE



- **Optimize for Efficiency:** Prioritize hardware that balances performance with energy efficiency. Generative AI models can be resource-intensive, so consider power consumption and explore options like specialized AI accelerators or cloud-based solutions.
- **Factor in Scalability:** If broader application is a possibility, choose hardware solutions that can scale easily to accommodate increased workload demands in the future. Consider cloud-based options with flexible resource allocation.
- **Evaluate Vendor Support:** When selecting hardware, factor in the quality of vendor support offered. Reliable technical support can be crucial for troubleshooting hardware-related issues during the pilot.



-
- **Overspend on Unnecessary Power:** Don't overshoot your needs with overly powerful hardware.
 - **Neglect Maintenance Costs:** Factor in ongoing maintenance costs associated with hardware solutions (e.g., cooling systems).
 - **Lock Yourself In:** Avoid vendor lock-in by ensuring chosen hardware allows some degree of flexibility with different software options.



9

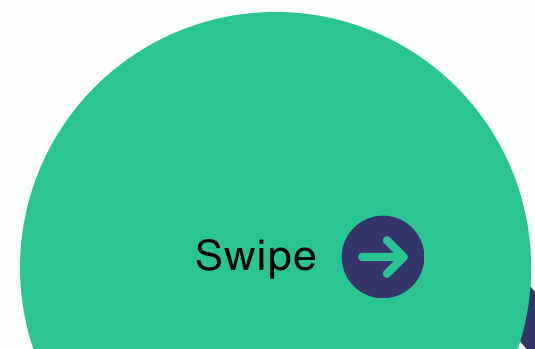
CHOOSE THE RIGHT SOFTWARE




- **Match Software to Deployment Approach:** Align your software choices with your deployment approach.
- **Prioritize User-Friendliness (if applicable):** Prioritize user-friendly software for non-technical users.
- **Embrace Open Source (Consider):** Explore open-source software options, considering potential expertise needs.



- **Overlook Security Vulnerabilities:** Don't choose software with known security vulnerabilities.
- **Skimp on Licensing Costs:** Factor in long-term licensing costs and potential compatibility issues with future upgrades.
- **Neglect Compatibility:** Ensure chosen software is compatible with your existing hardware and data infrastructure.





10

IMPLEMENTATION: BUILD | INTEGRATE | TEST, REFINE



- **Build (MVP Focus):** Begin with a Minimum Viable Product (MVP). This streamlined version of your AI solution allows for faster testing and iteration, enabling you to learn and refine more quickly.
- **Integrate (Depending on the**
- **Test (Real-World Users):** Test your MVP with real users, whether internal teams or customers. This provides valuable insights into real-world usability and potential impact on the KPIs you defined.
- **Integrate:** Meticulously plan how the AI solution will integrate with existing systems and workflows. Consider data exchange formats, user experience within existing interfaces, and potential API integrations.
- **Refine (Embrace Iteration):** Anticipate that the initial MVP will require refinement. Use test results to iterate on your design, improve performance, and ensure alignment with your value hypothesis.



-
- **Build (Over-engineer):** Don't get bogged down building a complex system from the start. Prioritize a simple, functional MVP that can be tested and improved upon quickly.
 - **Integrate:** Don't Treat integration as an afterthought. A clunky integration process or disruption to existing workflows can hinder user adoption and limit the effectiveness of your AI solution.
 - **Test (Ignore KPIs):** Testing isn't just about functionality. Continuously monitor the KPIs established for each use case. Their performance reveals areas for improvement and the effectiveness of your generative AI solution.
 - **Refine (Neglect Documentation):** Document your decisions and learnings throughout the pilot. This documentation becomes a valuable resource for future iterations and scaling the solution.





11

MAKE A DECISION: STOP OR DEPLOY


- **Analyze Performance:** Evaluate the performance of each use case against the KPIs established in the value hypothesis.
- **Consider Options:** Based on the evaluation, explore different options:
 - **Stop:** Don't be afraid to stop a use case if it consistently underperforms or fails to deliver any meaningful impact on your KPIs. Cut your losses and focus on more promising use cases.
 - **Refine:** If a use case shows promise but requires further development, prioritize refining the model architecture, user adoption efforts, or data quality based on the identified challenges.
 - **Deploy:** If a use case demonstrates significant progress towards meeting your value hypothesis, consider deploying it up for broader adoption within your organization.



-
- **Fall Victim to Sunk Cost Fallacy:** Just because you've invested time and resources into a use case doesn't mean you have to continue if it's not delivering results. Be objective and prioritize use cases with the highest potential for success.
 - **Neglect User Feedback:** User feedback is crucial during this decision phase. Understand why users may not be adopting the solution and address those concerns before scaling.
 - **Ignore Data Insights:** Don't base your decision solely on gut feeling. Leverage the data collected throughout the pilot to identify specific areas for improvement or potential roadblocks that might be addressed through refinement.



Swipe




12

IF PILOT IS SUCCESSFUL: BUSINESS READINESS PREPARATION



- Business readiness is crucial for a seamless transition and maximizing the value of your AI exploration.
- **Business Readiness Check List (not Exhaustive):**
 - Integrations
 - Test Scenarios
 - Compliance and Security
 - Version Control
 - Back-ups
 - Documentation
 - User Training
 - Communication Plan
- **Other Considerations:**
- **Phased Rollout:** Explore a phased rollout approach, starting with a limited group of users or a specific department before scaling to a wider audience. This allows for controlled testing, quicker issue identification, and user feedback gathering before a broader launch.
- **Rollback Plan:** Develop a rollback plan in case of unforeseen issues after launch. This plan outlines the steps to revert back to a previous version of the AI solution or disable it entirely if necessary.



13

IF PILOT IS SUCCESSFUL: DEPLOYMENT



- **Monitor Closely:** Continuously monitor the performance of the AI solution once it's live. Track key metrics like model accuracy, system performance, user adoption rates, and error logs.
- **Communicate Proactively:** Keep stakeholders informed throughout the launch day. Provide updates on progress, address any initial concerns promptly.
- **Offer User Support:** Be prepared to offer user support on launch day. Have a dedicated support channel available to address user questions and troubleshoot any initial issues.
- **Gather User Feedback:** Actively solicit user feedback on launch day.



-
- **Ignore Unexpected Results:** Don't be afraid to adapt if you encounter unexpected results on launch day. Be prepared to make adjustments or troubleshoot based on real-time data and user feedback.
 - **Neglect User Onboarding:** Don't assume users will know how to use the AI solution immediately. Provide ongoing onboarding support, including access to training materials and quick reference guides.
 - **Underestimate Communication Needs:** Don't underestimate the importance of clear communication. Anticipate user questions and concerns,



CELEBRATE THE WIN
LEARN AND ADAPT

www.ikigai-edge.com

Swipe 



14


IF PILOT IS SUCCESSFUL: TRAINING & SUPPORT



- **Offer Diverse Training:** Cater to different learning styles with written guides, video tutorials, and workshops.
- **Dedicated Support Channels:** Establish phone, email, or chat support for user questions and troubleshooting.
- **Gather Continuous Feedback:** Proactively collect user feedback to improve training and support.



- **Skip User Onboarding:** Provide a structured onboarding process with training and dedicated support.
- **Generic Training:** Tailor training materials to user roles, skills, and specific use cases.
- **Ignore User Feedback:** Analyze user feedback to improve training materials and support channels.





15


IF PILOT IS SUCCESSFUL: MONITORING & OPTIMIZATION



- **Track Key Metrics:** Monitor model accuracy, system performance, and user adoption rates.
- **Analyze User Feedback:** Actively solicit user feedback to identify areas for improvement.
- **Embrace Continuous Learning:** Use data and feedback to iterate and optimize the AI solution.



- **Set & Forget:** Don't simply launch the pilot and forget it. Continuously monitor performance.
- **Ignore Feedback:** Don't disregard user feedback - it holds valuable insights for optimization.
- **Stagnant Approach:** Don't be afraid to adapt the AI solution based on what you learn.





16


IF PILOT IS SUCCESSFUL: SCALE & INNOVATE



- **Be Prepare for Scalability:** Ensure solutions (hardware, software) can handle wider adoption.
- **Refine for Efficiency:** Optimize the model and user interface based on learnings.
- **Invest in User Adoption:** Develop a strategy for user training and ongoing support as you scale.
- **Measure Success:** Define and track relevant metrics to measure the ongoing success of your scaled AI solution.
- **Build a culture of adaptability and collaboration**



-
- **Rush Scaling:** Scale thoughtfully, prioritizing solutions that can handle increased workloads.
 - **Neglect User Needs:** Don't forget about user experience when scaling - address potential challenges.
 - **Stagnation:** Don't stop innovating! Use learnings to explore new use cases and applications.
 - **Lack of Collaboration:** Don't allow information silos to develop.



APPENDIX

BUSINESS CASES EXAMPLES

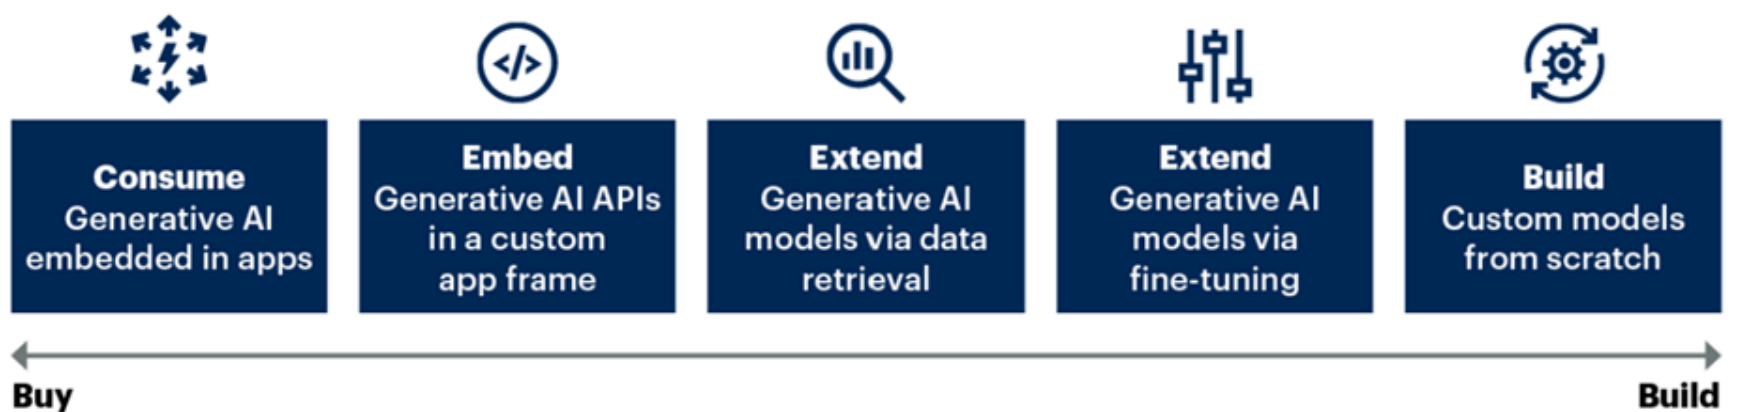
Department	Business Use Case	Impact
Sales	<ul style="list-style-type: none"> * Personalized Cold Email Outreach: Generate personalized cold emails that introduce the salesperson and highlight the prospect's specific pain points and how your product or service can address them. * Dynamic Sales Call Scripts & Follow-Ups: Create tailored call scripts leading to productive sales conversations. 	Increased lead engagement, improved conversion rates for qualified leads, more efficient use of sales rep time.
Marketing	<ul style="list-style-type: none"> * Automated Blog Content Creation: Generate blog post outlines, draft content that resonate with the target audience. * Creative Ad Copywriting: Produce variations of ad copy for different marketing channels to identify the most effective messaging. 	Increased content output, improved content quality and relevance, data-driven marketing decisions.
R&D	<ul style="list-style-type: none"> * Material & Molecule Generation: Utilize AI to design and generate materials with specific desired properties (e.g., strength, conductivity, biocompatibility) for applications in various industries. * Product Design Exploration & Prototyping: Leverage generative AI to explore a vast design space, generating and iterating on various product concepts based on defined parameters and user 	Faster material discovery and development, reduced time-to-market for new products, improved product innovation and differentiation.
Operations	<ul style="list-style-type: none"> * Demand Forecasting: Generate forecasts for product demand based on historical sales data, market trends, optimizing inventory management and production planning. * Predictive Maintenance in Manufacturing: AI can predict when equipment might fail and suggest solutions to avoid downtime. 	Reduced downtime, improved operational efficiency, lower maintenance costs, optimized inventory levels.
Finance	<ul style="list-style-type: none"> * Financial Report Summarization with Insights: AI can summarize financial reports and generate forecasts to help businesses make better decisions. * Scenario Modeling & Driver-Based Forecasting: AI can create automated scenario models that simulate the impact of various market changes, economic conditions, or competitor actions on your financial performance. 	Improved financial reporting efficiency, deeper insights, reduced workload for finance staff Improved financial planning accuracy, data-driven decision making for strategic resource
HR	<ul style="list-style-type: none"> * Personalized Job Descriptions: Generate targeted job descriptions that accurately reflect the required skills and experience for a specific role, attracting a more qualified pool of applicants. * Candidate Screening: Utilize AI to screen resumes and cover letters, identifying top candidates based on keywords, skills, and experience mentioned, streamlining the initial recruitment process. 	Streamlined recruitment process, improved candidate experience, reduced hiring time, identification of top talent.
IT	<ul style="list-style-type: none"> * Automated Incident Ticketing & Routing: Generate incident reports automatically based on system logs and alerts, routing them to the appropriate IT specialist based on the issue category and urgency. * AI-powered Chatbot Support: Develop a chatbot powered by generative AI to answer basic user questions about IT policies, troubleshoot common technical issues, and escalate complex problems 	Improved incident response time, reduced workload for IT staff, increased efficiency in resolving issues, improved user self-service capabilities.
Legal	<ul style="list-style-type: none"> * Automated Contract Drafting & Review: Generate first drafts of legal contracts based on pre-defined templates and specific client needs. 	Increased efficiency in contract creation, reduced time spent on repetitive tasks, improved risk identification and mitigation.
Project Management	<ul style="list-style-type: none"> * Automated Risk Identification & Mitigation Planning: Analyze project data (task dependencies, resource allocation, budget constraints) and generate reports highlighting potential risks and proposing mitigation strategies to proactively address them. * Smart Task Delegation & Scheduling: Utilize AI to analyze team member skills, workload, and project deadlines to suggest optimal task assignments and generate dynamic project schedules that maximize team efficiency and meet project milestones. 	Improved risk management, proactive problem-solving, reduced project delays, optimized team resource allocation.
Customer Service	<ul style="list-style-type: none"> * Personalized Chatbots & Virtual Assistants: Develop AI-powered chatbots and virtual assistants trained to answer frequently asked questions, troubleshoot common issues, and provide basic customer support. 	Reduced workload for human agents, faster resolution times for simple inquiries, improved customer satisfaction through 24/7 availability and personalization.



APPENDIX


DECIDE ON THE DEPLOYMENT APPROACH

Generative AI Deployment Approaches



Source: Gartner
797246_C

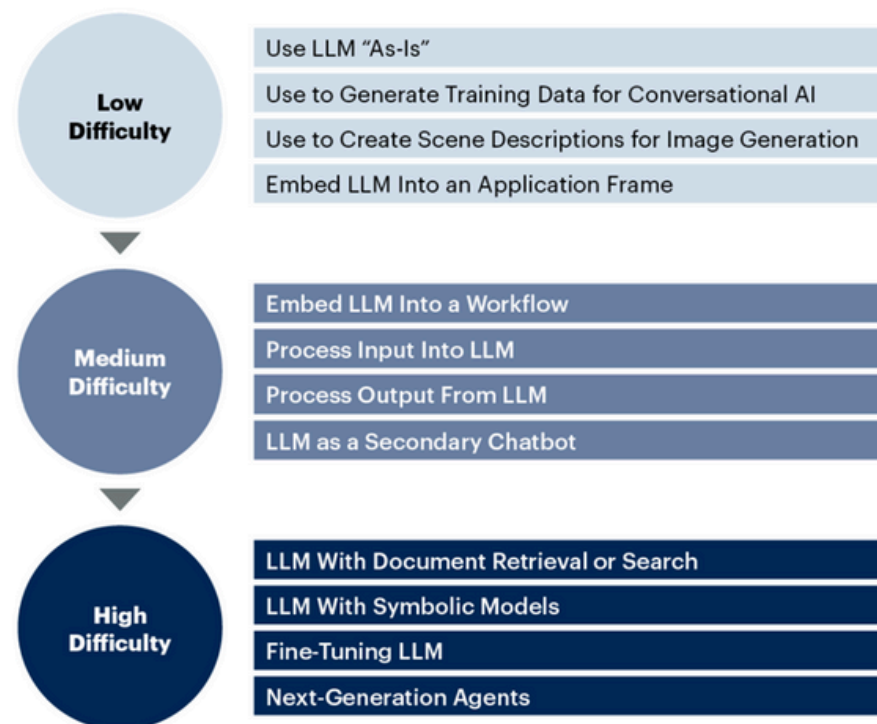
- **Consume:** Access pre-trained generative AI models offered by cloud providers or specialized platforms. Think "renting" AI capabilities.
- **Embed:** Integrate pre-trained models directly into your existing applications through APIs. Think "plugging in" AI functionality.
- **Consume (Extend via Data Retrieval):** While the core model is pre-trained, you can potentially improve its performance for your specific use case by feeding it with additional relevant data. Think of it like: "Renting" a car, but adding GPS data specific to your destination for a more optimized experience.
- **Embed (Extend via Fine-Tuning):** Fine-tuning allows you to further tailor a pre-trained model to your specific task or domain. This involves adjusting the model's internal parameters based on your own dataset. Think of it like: "Plugging in" a pre-built bookshelf, but then adding additional shelves or customizing the layout to better suit your book collection.
- **Build:** Develop and train your own custom generative AI models from scratch. Think "building your own" AI solution



APPENDIX

DECIDE PATTERNS

Design Patterns for Large Language Models (LLMs)



Source: Gartner
797246_C

Generative AI Design Patterns:

- Consume (LLM "As-Is"): Use a pre-trained model directly from a cloud provider, simplest option.
- Consume (Enhanced): Combine the pre-trained model with document retrieval for better results.
- Embed (App): Integrate the model into your existing application's user interface.
- Embed (Workflow): Automate tasks or provide decision support within a specific workflow.
- Build (Fine-Tuning): Customize a pre-trained model for your specific use case with your own data.
- Build (Agents): Create intelligent virtual assistants or chatbots powered by generative AI.
- Build (Advanced): Combine generative AI with rule-based systems for tasks requiring logic and creativity.

Choosing the Best Pattern:

- Deployment Approach: Consider how you'll use the AI (cloud-based, integrated, etc.).
- Use Case Complexity: Simpler tasks might not require complex builds.
- In-House Expertise: Choose a pattern that aligns with your team's skills.



IKIGAI EDGE
CONSULTING

Connect With Us

Gain practical insights to achieve your essential business goals with our expert guidance. We specialize in navigating the complexities of introducing and deploying Generative AI within your organization.



cmejia@ikigai-edge.com



WWW.IKIGAI-EDGE.COM

