ANTHROP\C

Planning to production: Best practices for implementing AI

Successful implementation of AI is iterative. Enterprises that are leading the way in AI transformation start small, evaluate thoroughly, and scale gradually before achieving their best results. With any new technology like Claude, optimal results require careful planning and ongoing refinement.

By choosing Claude, you're starting with a framework and foundation that are grounded in frontier-defining models, reinforced by Anthopic's commitment to safety and security.

Choosing the Right Claude Model: Capabilities, Cost, and Safety

Selecting the appropriate Claude model is crucial for the success and safety of your AI implementation. Here's what you need to know:

Key Considerations

Capabilities: Match the model approach to your task complexity. For example, our most capable Sonnet models can handle everything from nuanced reasoning to complex problem-solving in an enterprise environment. Sonnet is particularly strong for agentic coding, tool use, and customer facing agents. Haiku is ideal for straightforward classification or analysis particularly when cost is a consideration.

Cost and Speed: Balance your budget with performance needs. Start by assessing your task complexity and volume while also considering your budget and performance needs – higher capability models may actually be more cost-effective for complex tasks. As with all projects you should start small, evaluating the performance of your chosen model, and then scale based on your evaluation tests.

Safety and Security: Prioritize models with strong safety measures. All Claude models are designed to be helpful, honest, and harmless. Anthropic's leadership in AI research and focus on safety helps protect your brand and mitigate risks. Our models reduce the risk of harmful or biased outputs, position you for compliance with current and future AI regulations, and provide strong safeguards against data leakage or misuse.

Best Practices for Implementing Claude in Your Business

Successful Gen AI implementations follow a structured approach:

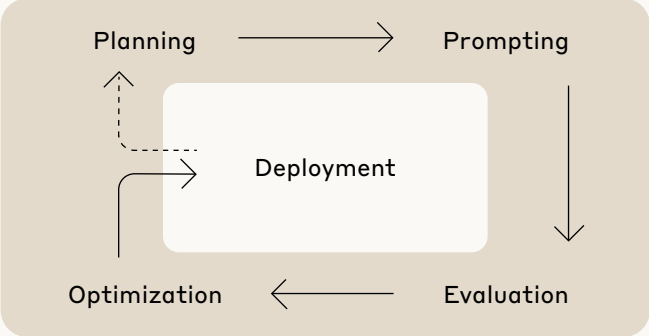

Let's dive into what that looks like in practice.

Planning

Start by selecting the right use case and model, and defining clear success criteria.

The right use case to start with should be well suited to current generative AI capabilities, such as understanding text and sentiment, classifying intent, generating content, or solving a repetitive problem or task. Then select the appropriate Claude model depending on your required combination of intelligence, speed, and cost. Make sure to establish specific, measurable success metrics that are

aligned with your business objectives. You need to be able to show a clear and quantifiable return on investment.

Many organizations are interested in building agentic systems as part of their AI implementations. These systems range from simple workflows with predefined paths to autonomous agents that dynamically direct their own processes. To succeed, you need to understand when agents are appropriate and how to implement them. Anthropic recommends starting with the simplest solution possible and only increasing complexity when needed. Optimizing single LLM calls with retrieval and in-context examples is sufficient for many applications. More complex implementations should be used only when they demonstrably improve outcomes. For detailed guidance on building effective agents, visit Anthropic's page on agents here.

Prompting

Good prompting, or prompt engineering, is simply adopting a strategic approach to communicating with AI systems to get the desired business outcomes. It's the difference between getting generic, unreliable outputs and achieving consistent, business-specific results that align with your organization's needs – whether that's generating analysis, writing code, or developing content.

At its core, a good prompt will provide a detailed task description and rules for how you want the model to handle it. Structure your prompts with clear task context, instructions, and desired output format. Include relevant background data and examples to guide the model.

Evaluation

You'll need to implement robust evaluations to assess model performance accurately. Start by creating a comprehensive evaluation dataset, including edge cases. Consider leveraging larger models as evaluators for smaller ones. Claude Sonnet 3.7 can augment human review when evaluating smaller models like Haiku, providing high-quality automated assessment at scale. For evaluating larger models like Sonnet itself, decompose the evaluation into smaller, more nuanced metrics. The most robust evaluation systems typically combine multiple approaches: rule-based evaluations implemented in code, LLM-powered evaluations using model evaluators, and targeted human grading for critical aspects.

Optimization

The results of your evaluation tests will help guide iterative improvements. Start by refining your prompts based on evaluation insights. Experiment with additional techniques like Chain of Thought (COT) reasoning, which encourages Claude to break down problems step-by-step, generating more accurate and nuanced outputs, and potentially highlighting where prompts are unclear. Few-shot prompting i.e. providing examples of correct answers to demonstrate desired outputs, is also a useful optimization technique.

Deployment

Roll out your Claude implementation strategically and prepare for ongoing refinement. Start with a small-scale pilot to identify and address issues early. Set up infrastructure for A/B testing of different prompts or models. Develop UI elements for human feedback and oversight, allowing users to interact directly with the system to flag concerning or incorrect outputs, label edge cases, and check results. Evaluation should be ongoing–you'll need to regularly update your offline evaluations based on production data.

TIPS FOR THE ROAD TO IMPLEMENTATION

Team Up with Experts: Leverage Anthropic' and our partners for a range of supports–from custom evaluations to deployment strategies.

Stay Current: Keep an eye on Claude's evolving capabilities to maximize your implementation's potential, including regularly reviewing documentation, guides, and code samples at docs.anthropic.com.

Prepare for the Future: Progressing with implementation is not a one-time task. It takes time to iterate and improve. Anthropic is at the forefront of AI development, consistently pushing the boundaries of what's possible. Partnering with Anthropic means you can build more confidently with frontier AI you can trust to be safer, more secure, and more reliable.

Visit docs.anthropic.com to explore our comprehensive developer documentation and cookbooks to get started building. You can also experience Claude's capabilities firsthand at console.anthropic.com, or connect with our Sales team to discuss your specific needs. The future of AI is evolving rapidly – let's build it together.