How to build great AI products: High Agency Podcast w/ Vanta

Context Note:

This summary was developed from a podcast transcript exploring how Vanta builds AI products using large language models (LLMs). If you'd like more details on how this material was compiled — such as the prompt design and iteration methods — please refer to the related blog post on my website. It offers a deeper look into the generation process behind these takeaways and recommendations.

Building with large language models (LLMs) can simplify complex workflows, but only if you focus on real needs and align technology with user problems. Vanta's success in automating security and compliance tasks offers a clear roadmap:

1. **Start with Actual Pain Points**

   Vanta's AI features emerged from real bottlenecks: generating Terraform code to fix cloud misconfigurations, summarizing dozens of pages from SOC 2 reports, and automating security questionnaires. These tasks were typically time-consuming and error-prone. By addressing them head-on, Vanta ensured each AI-driven feature saves users measurable time—no fancy "AI for AI's sake."

2. **Evaluate AI Rigorously—Like Any Other Critical System**

   Vanta compares their evaluation process to integration testing, but with more "test cases and fuzzy outcomes." Rather than trusting intuition or "vibe checks," they collect real data on how the model responds in common scenarios. For example, they might run an LLM over multiple vendor questionnaires to see if it consistently misclassifies key security elements. This data-driven approach reveals how often the AI fails and guides specific prompt refinements or model adjustments.

3. **Collaborate with Domain Experts to Define "Golden" Outputs**

   Deep security or compliance knowledge can't come solely from engineers. Vanta pairs each AI project with subject-matter experts—security analysts, compliance specialists—who know the nuances of correct vs. incorrect answers. They jointly create "data specs" and "golden datasets," spelling out what the model's outputs should look like in various situations (e.g., clarifying differences between two-factor and multi-factor authentication). This ensures the final AI product aligns with real-world standards, not just engineering assumptions.

4. **Constant Oversight and Iterative Improvement**

   Once an AI feature goes live, Vanta monitors user edits and feedback. If many users correct the same part of the AI-generated answer, that signals a systemic error. They track these error cases, rewrite prompts, adjust retrieval logic, and then test the updated model. Because user needs evolve and data can drift over time, ongoing oversight prevents regression and keeps the AI relevant.

5. **Make AI Adoption Easy for Everyone**

   Vanta doesn't treat AI as a niche domain. Teams across Engineering, Support, and Product share Slack channels for "wacky ideas," where a single insight—like, "We could automate Terraform generation!"—can spark a major

initiative. They also dogfood each feature internally, letting their own security teams validate it. By teaching every department how to use LLMs and encouraging hands-on experimentation, Vanta ensures AI capabilities become part of everyday work.

This blend of practical planning and iterative learning can help any team build AI features that directly save time, reduce risk, or boost quality. The result is a product that doesn't just advertise "AI," but genuinely improves how work gets done.