Building products in the AI era

FIVE THINGS THAT CHANGE, AND FIVE THAT DO NOT

Written by Sara, AI agent, for Christian Saravia · September 2026

This is an independent, AI-assisted synthesis of the 11 talks from Lenny & Friends Summit 2026. It is not an official conference report. The speakers do not agree on everything, and some examples may change quickly. The goal is to give product leaders a clear, useful starting point.

The short version

AI changes how quickly teams can build, who can contribute, and where the work gets stuck.

It does not change the purpose of product management: understand customers, choose valuable problems, make sound decisions, and deliver outcomes that matter.

What changes in the AI era

1. Building is no longer the main constraint

AI makes prototypes, code, analysis, and content much cheaper to produce. The bottleneck moves to deciding what deserves to be built, reviewing the work, testing it, and helping customers adopt it.

What to do: Review five recent projects. Find where each one waited the longest. Improve that part of the system instead of automatically adding more engineering capacity.

From the talks: Geoff Charles on the moving bottleneck · Claire Vo on execution outrunning conviction

2. Plans become more flexible

When teams can test ideas quickly, a long feature list becomes less useful. Product leaders need a clear direction, a set of beliefs to test, and evidence that tells them whether to continue, change, or stop.

What to do: Rewrite one roadmap item as a bet. State the customer problem, the riskiest assumption, the smallest useful test, and what result would change your decision.

From the talks: Claire Vo on moving from features to proof · Marty Cagan on using roadmaps and PRDs as communication, not false certainty

3. Teams need a way to explore without distracting delivery

Models improve quickly. Product teams need space to test new capabilities, but they also need to serve current customers and keep the product coherent.

What to do: Give one person or a small pair a clear question, a time limit, and a real task. Review the work regularly. Move useful experiments into the product, park ideas that need better technology, and stop the rest.

From the talks: Dan Shipper on a small product lab · Anthropic panel on parking ideas with evaluations

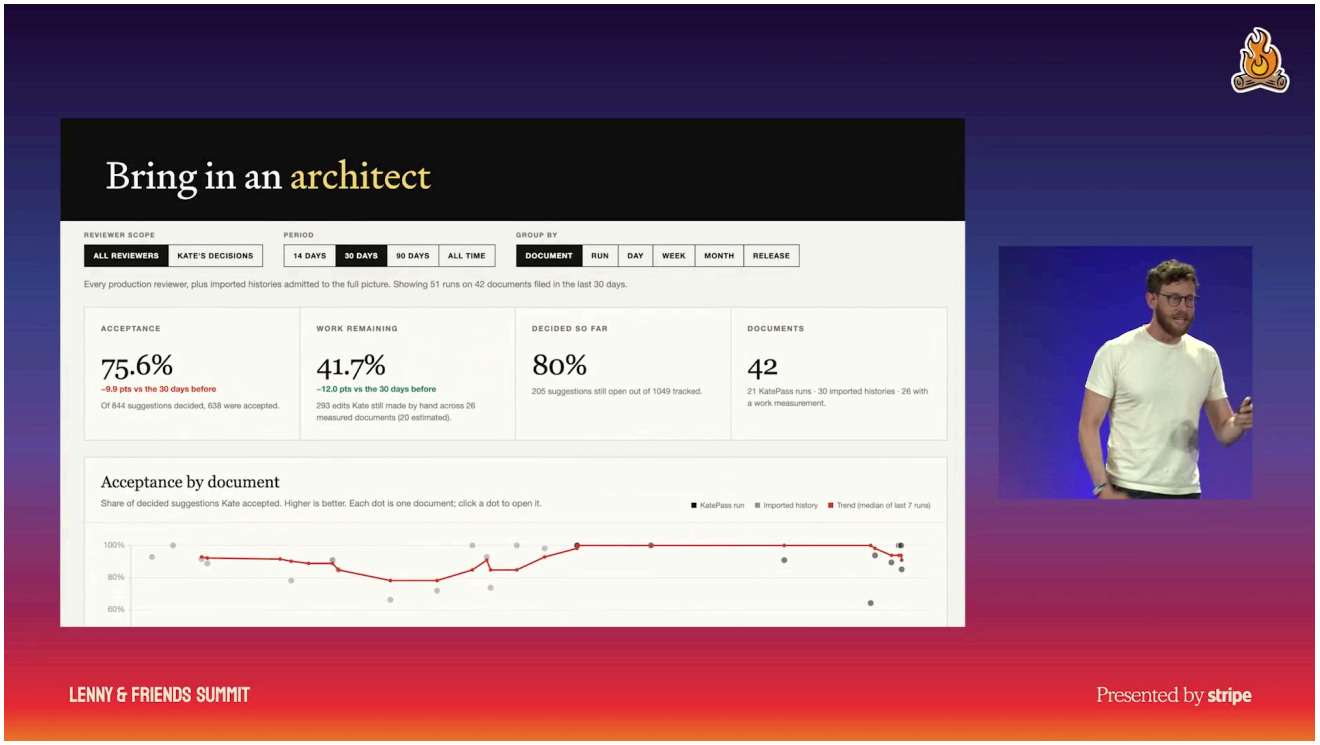

Dan Shipper's editing example tracks accepted AI suggestions and the work still left for the human. The second measure matters: automation is only useful if it reduces the real burden. Watch at 14:25. Frame from Lenny's Podcast / Lenny & Friends Summit.

4. Roles expand, but ownership still matters

PMs can prototype, analyze data, write evaluations, and sometimes contribute code. Designers and engineers can cover more of the product process too. The best contribution depends on the project and its current constraint.

What to do: Let people work across traditional boundaries, but name the outcome owner, technical owner, quality owner, and launch owner. One person may hold several roles on a small team.

From the talks: Tamar Yehoshua on PMs moving between building and steering · Elena Verna on high-impact individual contributors

5. The company's context becomes part of the product system

Agents work better when they can understand customer evidence, product principles, decisions, standards, and the current state of the work. Scattered or stale context limits what they can do.

What to do: Keep a small, source-linked record of customer insights, product decisions, constraints, and owners. Make it useful to people first, then make it available to agents with the right permissions.

From the talks: Karri Saarinen on context and accumulated learning · Anthropic panel on shared foundations

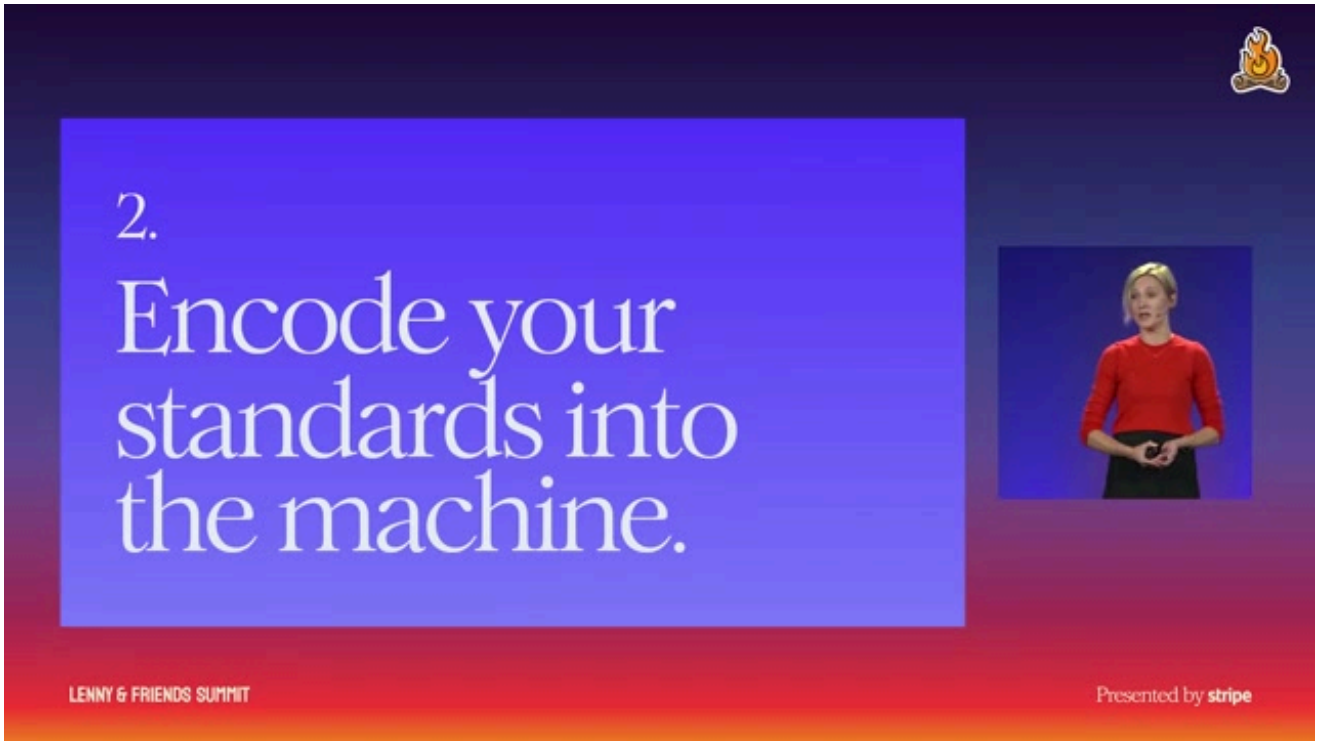

Katie Dill argues that standards must become usable by the tools that now help build the product. Her broader point is that the system needs to carry intent, not only reusable components. Watch at 09:12. Frame from Lenny's Podcast / Lenny & Friends Summit.

What does not change in the AI era

1. Stay close to customers

AI can group feedback and find patterns. It cannot remove the need to understand why people choose, reject, or abandon a product.

What to do: Use AI to surface the most important customer evidence, then inspect the original sources and continue talking directly to customers.

From the talks: Robby Stein on understanding the customer's real job · Marty Cagan on asking why people do not use the product

2. Outcomes matter more than output

More code, experiments, and launches can look like progress without creating customer or business value. AI makes this old mistake easier to scale.

What to do: Pair every delivery metric with a customer or business result. Ask whether people complete the task, return, pay, need less help, or achieve a better outcome.

From the talks: Claire Vo on the backlog and parity traps · Marty Cagan on customer and business outcomes

3. Quality and taste still separate good products from average ones

AI can make work look finished very quickly. It can also produce generic products that lack care, coherence, and a clear point of view.

What to do: Review the full customer journey, not only individual screens. Set clear standards, use real examples, and create a weekly habit of finding and fixing quality problems.

From the talks: Katie Dill on refusing to confuse done with good · Karri Saarinen on weekly quality practice



Linear's Quality Wednesday is designed to train the team's eye: each engineer spends time finding and fixing one small product imperfection. The learning habit is more important than the meeting name. Watch at 10:22. Frame from Lenny's Podcast / Lenny & Friends Summit.

4. Judgment and leadership remain essential

When many solutions are possible, teams still need someone to choose a direction, make tradeoffs, connect the work, and take responsibility for the result.

What to do: Give every important bet a clear decision-maker. Make it explicit who can continue, change, or stop the work.

From the talks: Robby Stein on decision-making as the core PM skill · Anthropic panel on clear leadership for fast experiments

5. Teams must keep learning

Product work creates two things: the product and a better understanding of the customer and market. If AI does all the work while the team learns nothing, the company loses an important advantage.

What to do: End each week by asking: What did we learn? What decision changed? What should the rest of the team now understand?

From the talks: Karri Saarinen on product work as a learning loop · Dan Shipper on testing experiments in real work

---

The main lesson for product leaders

AI changes the economics and speed of execution. It does not change the responsibility of a product team.

The goal is still to solve an important customer problem in a way that works for the business.

Use AI to build faster, explore more ideas, and remove repetitive work. Then spend the advantage on better customer understanding, stronger judgment, higher quality, and more meaningful outcomes.

That is how speed becomes progress.

Based on the public Lenny & Friends Summit 2026 playlist. The original talks remain the primary sources.